LA-UR-22-24475
Approved for public release; distribution is unlimited.

Title:	ASCR Reverse Site Visit
Author(s):	Qualters, Irene; Hagberg, Aric Arild; Shipman, Galen M.; Chacon, Luis; McCormick, Patrick Sean; Sornborger, Andrew Tyler; Iyer, Srinivas
Intended for:	Report
Issued:	2022-05-12

Los Alamos
NATIONAL LABORATORY

ASCR Reverse Site Visit

Irene Qualters, Associate Lab Director for Simulation and Computation (ALDSC), ASCR Executive

Srinivas Iyer, Office of Science Program Manager

Aric Hagberg, CRCL representative, Deputy Division Leader

Luis Chacon, Pat McCormick, Galen Shipman, ASCR PMs

Andrew Sornborger, LANL PI for QSC

May 12, 2022

Los Alamos delivers national security solutions

- **We are dedicated to addressing complex national security issues and the world's most difficult challenges**
  - By advancing and applying multidisciplinary science, technology & engineering capabilities;
  - In unique experimental, computational, and nuclear facilities;
  - With an agile, responsive, and innovative workforce;
  - And by partnering with peer institutions for mission success

LANL Statistics

- $4B budget
- 40 square miles, 47 technical areas
- 897 bldgs., 8.4M sq ft,
- 13 nuclear facilities
- > 13,800 workers
- 9,800 career employees
- 1,637 students, 462 postdocs
- Employee average age: 43
- 65% male; 35% female 45% minorities
- 40.2% of employees are native New Mexicans

Our Lab has grown by more than $1B in two years

- LANL has grown by **$1.48B and ~2,500 staff** in the 3+ years since Triad
- Managers hired 1,277 regular/term employees in FY21
  - Reduced hiring cycle from 125 to 69 days
  - Poised to launch Lab-wide on-the-spot hiring
- Challenge now is to grow our human and infrastructure capacity so we can keep pace with budget increases

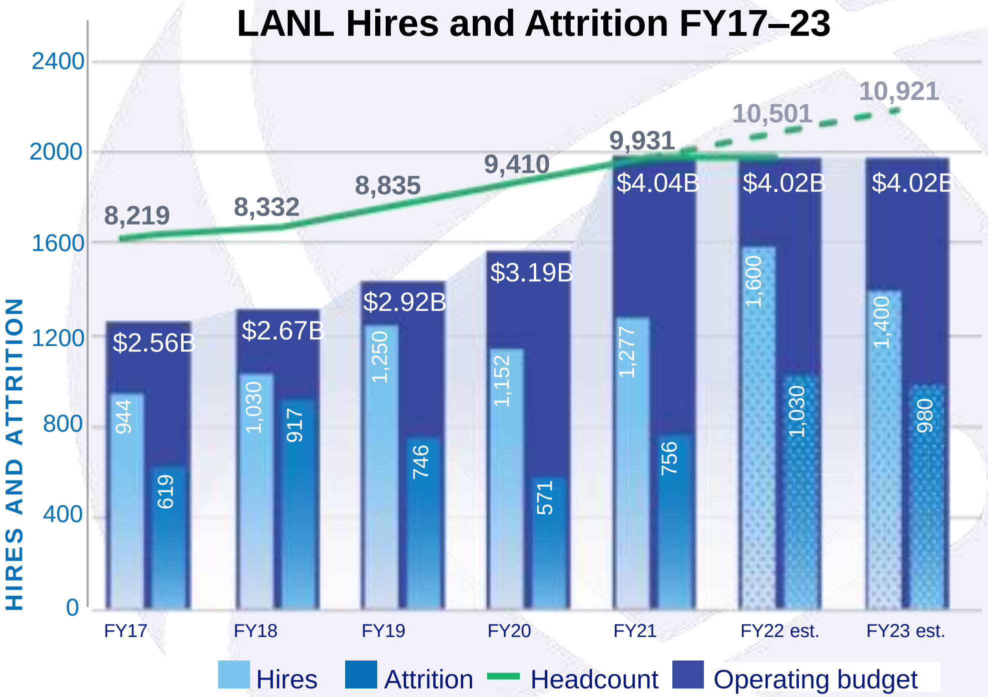

LANL Organization gives “line of sight” for DOE/OS engagement

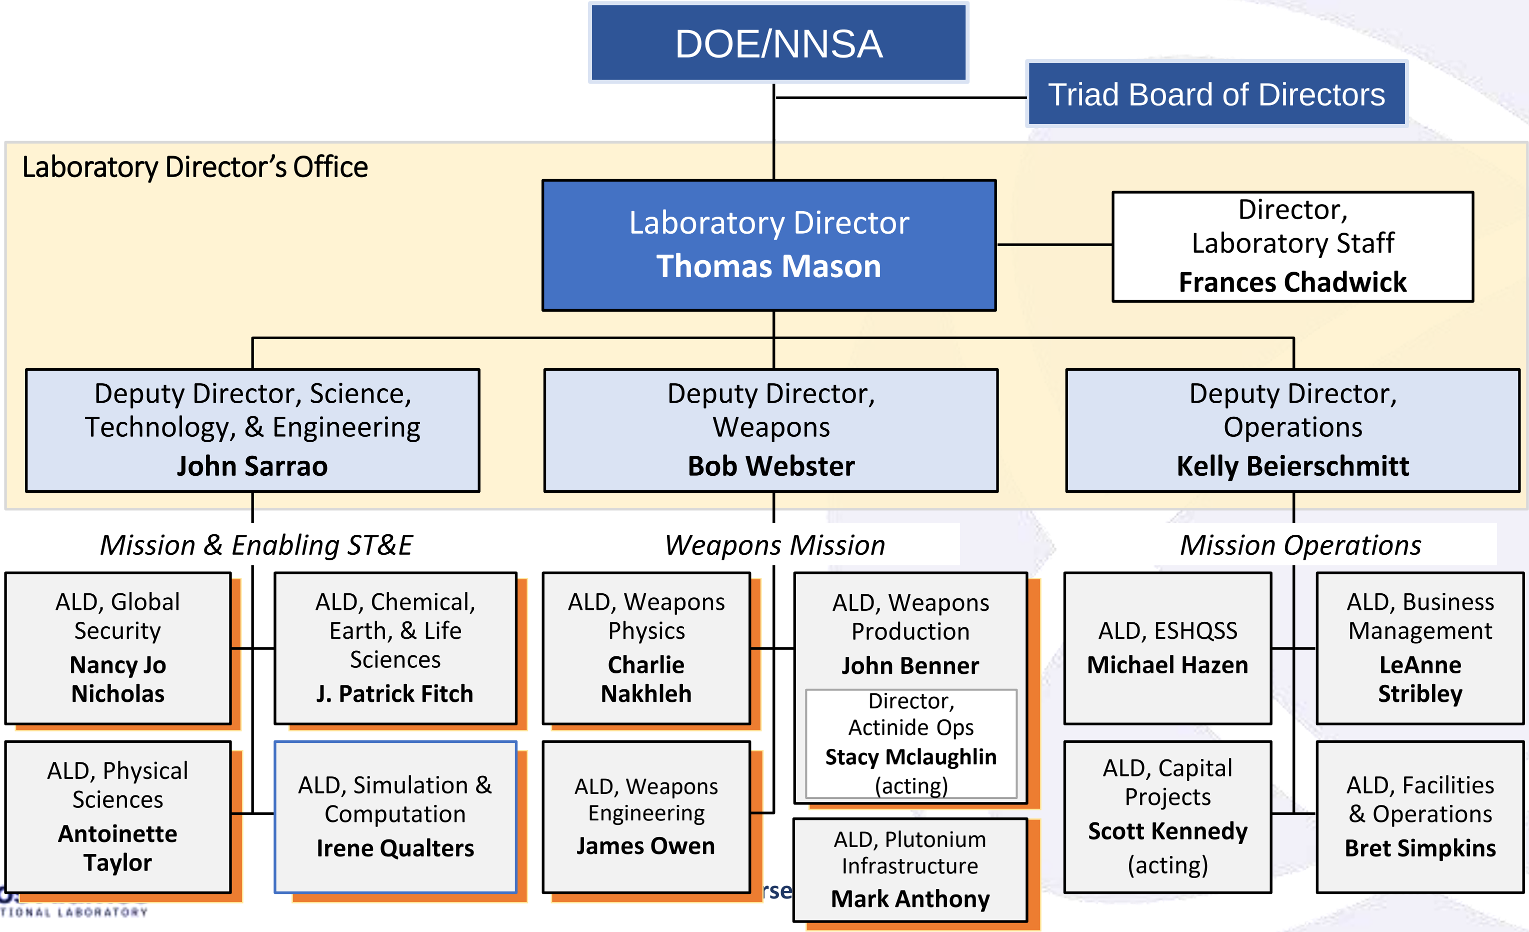

Customer Interface Plan for DOE Office of Science Programs at Los Alamos National Laboratory

Roles & Responsibilities

- **Deputy Lab Director for STE**
  - Serve as an ambassador and senior-level spokesperson for the Laboratory, including **point of contact with the SC Director**.
- **Associate Lab Directors**
  - Function as Lab's Senior Executive Points of Contact (in conjunction with assigned SC ADs) to develop, build, and maintain strong business and personal relationships.
- **SC Program Director**
  - Serve as the overall institutional relationship manager for DOE-SC by coordinating overall communication between the Laboratory and DOE-SC; Manage SC projects at the Laboratory as an institutional portfolio.
- **Lab Program Manager(s)**
  - Serve as Laboratory Points of Contact with assigned DOE-SC Division Leader(s) and their program managers to develop, build, and maintain strong business and personal relationships; Coordinate overall communication between the Laboratory and assigned SC Division(s).

Our Lab Agenda emphasizes Simultaneous Excellence

How **we do our work is as important as** ***what*** **we do**

Strategic Objectives
Nuclear Deterrent
Threat Reduction
Technical Leadership
Trustworthy Operations

Critical Outcomes
Pit Production
Non-Nuclear Production
Computational Breakthroughs
Experimental Advances
Integrated Deterrence
Technology Modernization
Threat Reduction
Quantum Leadership
Climate & Clean Energy
Biosecurity Preparedness
Culture Enhancements
Operational Capacity
Force for Good

Technical Leadership

Deliver scientific discoveries and technical breakthroughs to advance relevant research frontiers and anticipate emerging national security risks

Long-term ST&E stewardship is based on 6 Capability Pillars which define key areas of science, technology and engineering in which we must lead

ENGINEERING

MATERIALS FOR THE FUTURE	Defects and Interfaces Extreme Environments Emergent Phenomena
NUCLEAR AND PARTICLE FUTURES	Accelerator Science & Technology Applied Nuclear Science & Engineering High Energy Density Plasmas & Fluids Nuclear, Particle, Astrophysics & Cosmology
INTEGRATING INFORMATION, SCIENCE, AND TECHNOLOGY FOR PREDICTION	**Computing Platforms** **Computational Methods** **Data Science**
SCIENCE OF SIGNATURES	Nuclear Detonation Nuclear Processing, Movement, Weaponization Natural and Anthropogenic Phenomena
COMPLEX NATURAL AND ENGINEERED SYSTEMS	Human–Natural System Interactions: Nuclear Engineered Systems Human–Natural System Interactions: Non-Nuclear
WEAPONS SYSTEMS	Design Manufacturing Analysis

The LDRD components combine for a balanced investment in LDRD objectives

Exploratory Research
27%
Innovate at the frontiers of technical disciplines

Directed Research
45%
Create multidisciplinary solutions to complex problems defined by Lab strategy

Reserve (unencumbered)
4%

Director's Initiatives
7%
Invest in the Lab Agenda with the rigor and creativity of LDRD

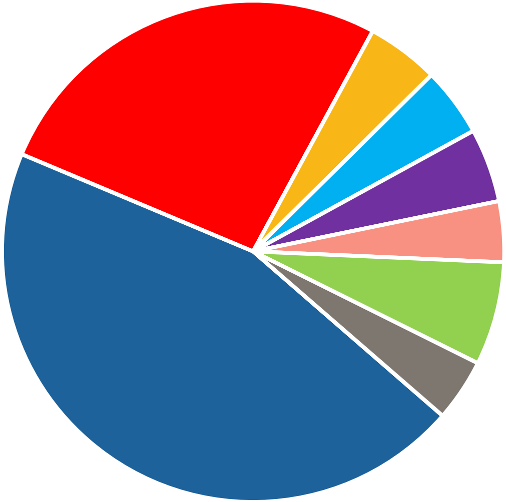

Baseline
LDRD Program
FY22 Budget: $180M

Early Career Research
5%
Develop next-generation technical leaders

Postdoctoral R&D
4%
Attract and recruit top-quality talent into the Lab's pipeline

Mission Foundations
5%
Translate discovery into novel mission solutions

Centers Research
4%
Incubate emerging ideas and talent in areas defined by the Lab's Strategic Centers

Focus on people

- Hybrid student program successfully realized in 2021; planned for 2022
- Summer schools are a unique pipeline to recruit students in key technical areas
- Rigorous postdoc conversion process leads to talented early career staff with key mission skills
- Student programs, pipeline initiatives boost diversity in student pipeline
  - See e.g., women.lanl.jobs

30.2% of regular/term employees have at least one degree from a NM college or university

Los Alamos NATIONAL LABORATORY

Student/Postdoc Numbers (FY16–21)

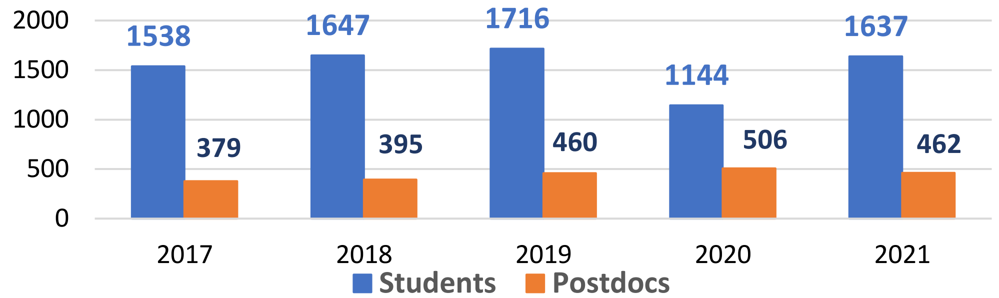

Postdoc Conversions compared to total non-management technical hires (FY17-FY22 Q1)

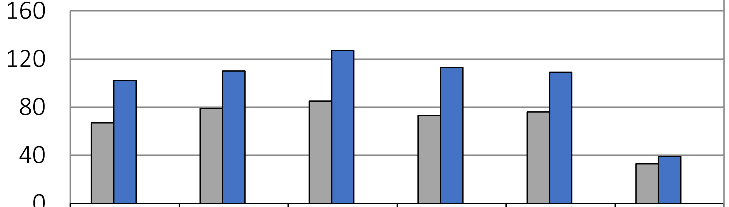

	FY17	FY18	FY19	FY20	FY21	FY22 thru Q1
Postdoc Conversions to Scientist and R &D Engineer	67	79	85	73	76	33
Total PhD Non-Management Scientist and R&D Engineer, Levels 1-3 plus PhD GRA & Postdoc Conversions	102	110	127	113	109	39
% of Postdoc Conversions to Total PhD Staff Hires	66%	72%	67%	65%	70%	85%

ALDSC Regular and Term attrition has outpaced external hires and conversions by 6 YTD

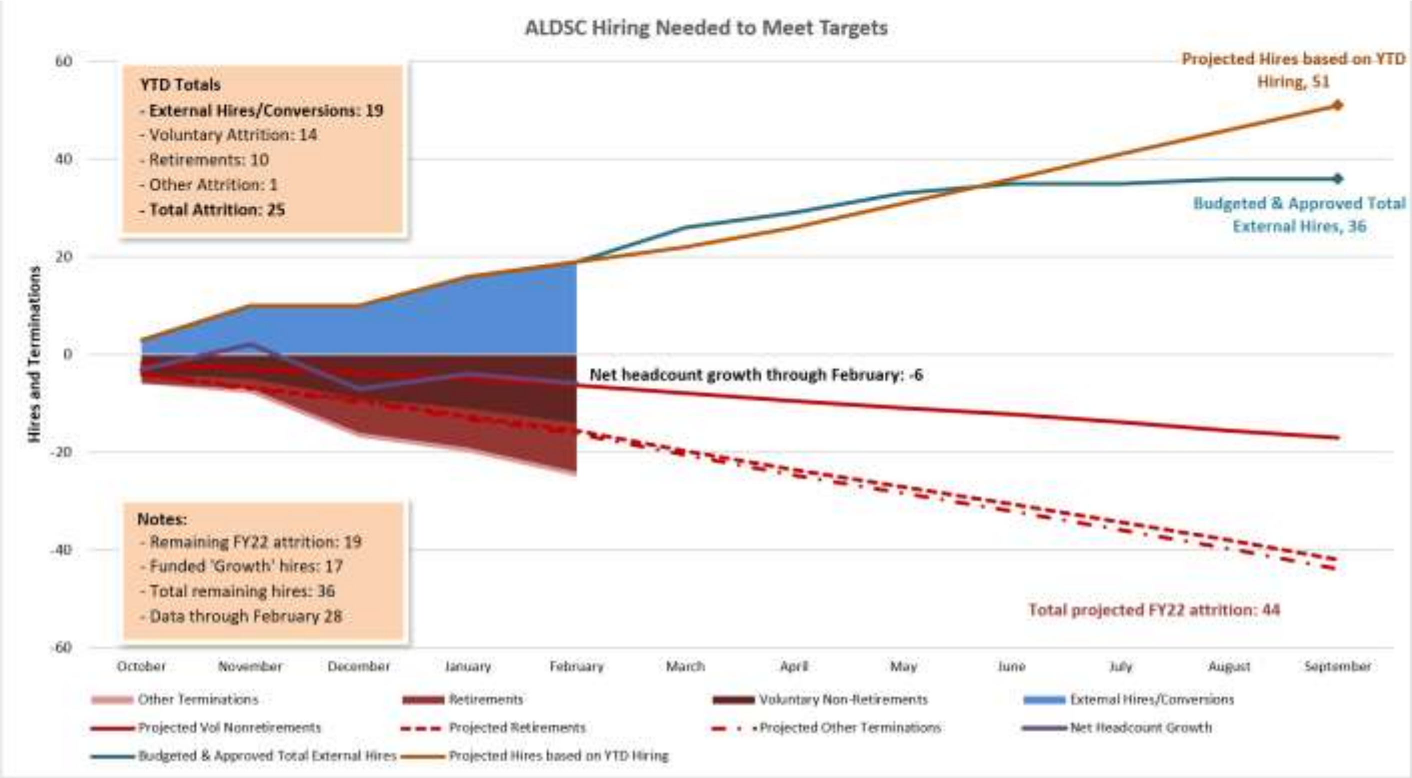

- February totals
  - External Hires & Conversions: 3
  - Terminations: 5
- The YTD pace projects 51 external Reg/Term hires/conversions.
- Total budgeted growth hires plus YTD hires projects 36 Reg/Term hires/conversions.

- Blue line shows budgeted external hires and conversions planned in Finance's Forecast system.
- Brown line shows the projected total hires based on YTD hires and historical hiring distribution.
- The gap between the blue and the brown line represents an underrun of committed funds.

Agenda

Srinivas Iyer

Irene Qualters

Aric Hagberg

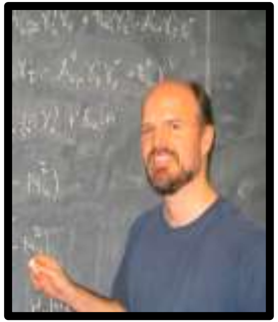

Andrew Sornborger

Luis Chacon

Pat McCormick

Galen Shipman

ASCR Portfolio: 1 hour

- **Introduction** (Irene Qualters)
- **Applied Mathematics Research** (Luis Chacon)
- **Computational and Data Partnerships (**Galen Shipman)
- **Computer Science/Workflows/Edge Research** (Pat McCormick)

Break 30 mins.

Crosscuts: 1 hour

- **Scientific Data** (Pat McCormick)
- **Quantum Information Science** (Andrew Sornborger)
- **Artificial Intelligence/Machine Learning** (Aric Hagberg)

6/8 LANL 2021 R&D 100 awards are highly resonant with ASCR

Demonstrate DOE/Labs excellence in Multidisciplinary R+D+Partnerships

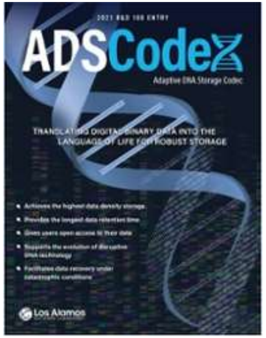

ADS Codex: Adaptive DNA Storage Codec

CICE Consortium

Quantum Ensured Defense of the Smart Electric Grid (QED)

Smart Tensors AI Platform

MOCHI

Terra Spotlight

In spite of strong post-doc and student pipelines, ALDSC is experiencing unprecedented attrition across early, mid and late-career staff

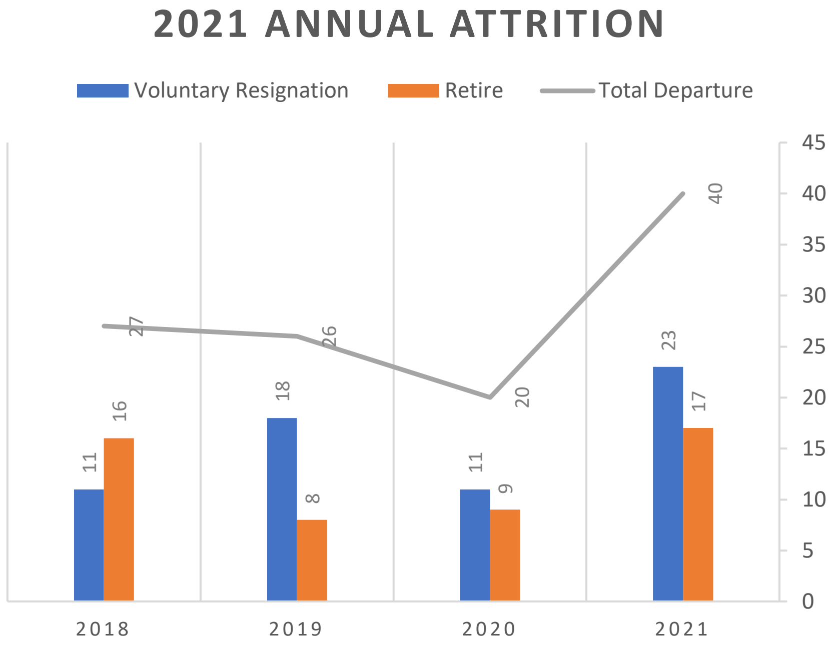

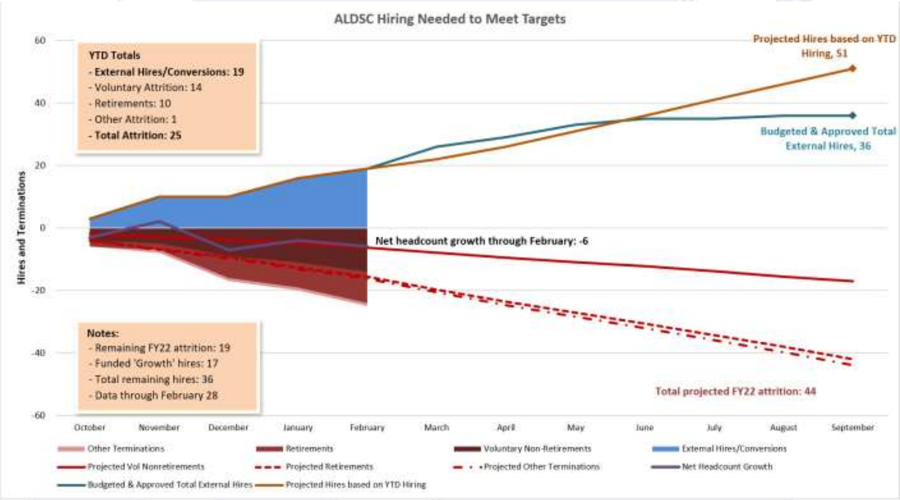

Overview of Personnel in ASCR Program Areas

- Total headcounts in ASCR program areas are estimates (> 250)
  - Many LANL programs engage math, cs, computational science researchers including large NNSA/ASC program and $150m/year LDRD program
  - ASC, SPP, and LDRD have R&D programs in ML and Quantum (with security focus)
  - Applied Math also includes approximately 25 research statisticians
- Students: Significant summer hiring (~100 in ASCR relevant summer schools, + more in internships). **We are executing a mix of in-person (primary) and remote student (secondary) summer programs for summer 2022.**
- As ECP trends towards Development, Deployment and Operations, we are encountering greater pull from industry without the offsetting competitive advantage of multidisciplinary R&D.

Overview of Personnel in terms of Job Type (FY21)

		ASCR					ECP			OTHER
		KJ0401	KJ0402	KJ0403	KJ0404	KJ0405	AD	ST	HI	
Early Career	Heads	2	2	11	0	0	12	6	4	35
	FTEs	0.59	0.89	2.84	0	0	4.77	2.81	0.8	
Staff	Heads	0	1	18	0	0	21	12	8	70
	FTEs	0.05	0.41	3.77	0	0	6.96	4.48	2.05	
Postdocs	Heads	0	0	12	0	0	8	1	0	20
	FTEs	0	0	3.18	0	0	2.58	0.17	0	
Students	Heads	0	5	7	0	0	9	2	1	10
	FTEs	0.02	1.42	2.21	0	0	2.53	0.66	0.03	
Support Staff	Heads	0	0	0	0	0	0	0	0	0
	FTEs	0	0	0	0	0	0	0	0	
Totals	Heads	2	8	48	0	0	50	21	13	
	FTEs	0.66	2.72	12.0	0	0	16.84	8.12	3.15	135

Overview of Personnel in ASCR Program Areas (FY21)

ASCR Program Area	Total Headcount	Funding Source (Headcount)						
		ASCR	ECP AD	ECP ST	ECP HI	NQISRCs (LBNL and ORNL centers only)	Other (includes LDRD)	
							Headcount	Funding sources
Applied Math (Total) KJ0401	37	2 Staff					~35	LDRD, NNSA
Computer Science (Total) KJ0402	34	3 Staff 3 Students		12 Staff 1 Postdoc 2 Students	12 Staff 1 Student		~20	LDRD, NNSA, Global Security
Computational Partnerships (Total) KJ0403	188	20 Staff 2 Postdocs 7 Students	33 Staff 8 Postdocs 9 Students			18 Staff 11 Postdocs	>80	LDRD
Advanced Computing Technologies (Total) KJ0503								
ECP only (Total)								
NQISRCs only (Total)								
Total	259	25 Staff 2 Postdocs 10 Students	33 Staff 8 Postdocs 9 Students	12 Staff 1 Postdoc 2 Students	12 Staff 1 Student	18 Staff 11 Postdocs	>135	

Overview of Personnel in terms of Job Type (FY21 Data)

Job Type	Total #Headcount / FTEs	Headcount and FTEs per funding category						
		ASCR Research	ECP AD	ECP ST	ECP HI	NQISRCs	Other	
							Headcount / FTEs	Funding source
Senior/Mid-Career Staff	133	9/2.3	21/6.96	12/4.48	8/2.05	13/2.22	70	NNSA, LDRD, Global Security
Early Career (< 10 years last degree)	73	11/3.48	12/4.77	6/2.81	4/0.8	5/0.89	35	NNSA, LDRD, Global Security
Post Docs	22	2/0.64	8/2.58	1/0.17	0/0.0	11/2.62	20	LDRD, NNSA
Undergraduate + Graduate Students	34	12/3.45	9/2.53	2/0.66	1/0.3	0/0.2	10	LDRD, NNSA
Support Staff	0	0/0	0/0	0/0	0/0	0/0		
Total	262	34/9.87	50/16.84	21/8.12	13/3.15	29/5.94	135	

Overview of ASCR Program Funding for next 5 years

(in thousands of $)

ASCR Program Area	FY2022	FY2023	FY2024	FY2025	FY2026
Applied Math (Total)	900	900	550		
AI/ML for Science	400	400	400		
Other	500	500	150		
Computer Science (Total)	606				
Scientific Data	276				
Other	330				
Computational Partnerships (Total)	2200	1946	615	615	
SciDAC	965	965	615	615	
Quantum Information Science	1235	981			
Advanced Computing Technologies (Total)	4820	4820	4820		
Quantum Internet Testbeds					
NQISRCs	4820	4820	4820		
ECP (Total)	12030	7592			
AD	6599	4469			
ST	4366	2917			
HI	1065	206			
Total	20556	15258	5985	615	

Applied Math

Luis Chacon

Applied Mathematics and Plasma Physics Group, Theoretical Division

LANL ASCR Applied Math POC

May 12th, 2022

Applied Math at a glance

- **Applied Math portfolio** is centered in the Applied Mathematics & Plasma Physics Group (T-5) and builds on long LANL tradition in mathematical methods and numerical analysis.
- **Total effort** ~ 2 FTE among 6 staff & 1 post-doc. Applied Math community at LANL much larger at 45+
- **Currently ASCR portfolio [$0.9M/yr]**
  1. Base program: Enabling Multiphysics/Multiscale Plasma Simulations by the Development of Stable, Accurate, and Scalable, Computational Formulations and Solution Methods – LANL PI: Chacon [FY21-23; $350k/yr]
  2. Data Intensive ML: Inertial neural surrogates for stable dynamical prediction – PI: Tang [Burby co-PI, FY 22-24; $400K/yr]
  3. ASCR Mark Kac PD Fellowship: Harnessing the power of quantum computers for dynamical simulation and machine learning – PI: Z. Holmes (1st Mark Kac Fellow, $150K/yr)
- **Mark Kac Applied Math Post-doc Fellowship**
  - Started in FY21; $150k/yr from ASCR, co-sponsored by Center for Nonlinear Studies @ LANL
- **Current LANL AM investment areas**
  - Scientific AI/ML**:** significant LDRD investment (~$15M/yr)
  - Multiscale algorithms ($4.5M/yr)
  - Traditional Applied Math themes: PDE discretizations, multiscale solvers, optimization, etc. ($2.2M/yr)

Mark Kac Applied Math Postdoc Fellowship Program

($150K/yr for 3 years)

Focusing on numerical analysis, solvers, multiscale mathematics, optimization, scientific data analysis, machine learning, and quantum computing algorithms

Co-sponsored by CNLS (50%), allowing for 1 PD/yr for 3 years.

An essential element of our staffing pipeline

Eugene Kikinzon
Advanced discretization of PDEs.
C++ libraries for polyhedral meshes.

Converted to LANL staff, 2018

Erik Skau
Sparse Representations and Machine Learning for Scientific Data Sets.
Inverse Problems.

Converted to LANL staff, 2020

Svetlana Tokareva
Advanced discretization of PDEs.
Scientific Machine Learning applied to PDEs.

Converted to LANL staff, 2019

Zoe Holmes
Quantum algorithms and Machine Learning

First Mark Kac Fellow

?

Kac Fellow Selected TBA May 18th
Reactive Hydrodynamics

CNLS and ISTI provide important access to visitor programs, workshops, summer student schools, and rapid response funding.

Barren plateaus preclude learning scramblers (Z. Holmes)

Scientific achievement

We prove a no-go theorem for learning an unknown scrambling process (entanglement spreading) with quantum machine learning (QML).

Significance and Impact

Our study places generic limits on QML when lacking prior information.

Research details

- We show that when learning a typical scrambling unitary (or typical random unitary), any variational ansatz is highly likely to have a barren plateau landscape, i.e., cost gradients that vanish exponentially in the system size.
- This implies that the required resources scale exponentially even when strategies to avoid such scaling (e.g., from ansatz-based barren plateaus or No-Free-Lunch theorems) are employed.

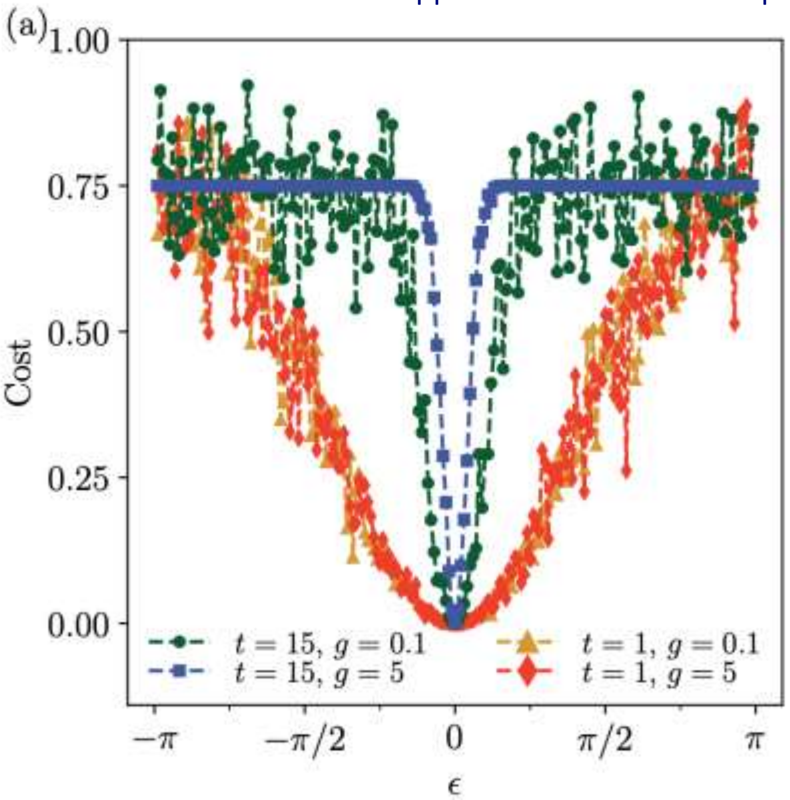

A cut through the landscape of a cost function used to learn an approximate scrambler. The parameters *g* and *t* indicate the entangling rate and scrambling time respectively. For large *g* and *t* the target is highly scrambling and the cost landscape forms a barren plateau with a narrow gorge in the region around the parameters that minimize the cost. In contrast, for a weaker scrambler the valley around the optimum is wider and the landscape is more featured.

Z Holmes, A Arrasmith, B Yan, PJ Coles, A. Albrecht, AT Sornborger, PRL **126** (19), 190501. (Work performed at LANL)

Enabling Predictive Plasma Simulations

(Lead: J. Shadid, SNL; LANL PI: Luis Chacon, $350K/yr)

LANL ASCR-funded research is at the forefront in multiscale simulation of inertial- and magnetic-confinement plasmas worldwide

- **Pioneered implicit solvers** and preconditioners for **plasma fluid hyperbolic models** (resistive MHD, extended MHD).
- **Pioneered fully implicit** algorithms for fully **kinetic particle (PIC) algorithms** while conserving charge and energy (**a 50-year-old quest**).
- **Pioneered** high-order/low-order (HOLO) **hierarchical multiscale solvers** for fully kinetic Eulerian and Lagrangian discretizations.
- **Pioneered fully implicit**, conservative **multiscale hybrid fluid-kinetic algorithms** by merging fluid and kinetic developments.
- Produced **first-of-a-kind physics contributions** in magnetic reconnection and magnetic fusion confinement systems.

LANL ASCR-funded research is having strong impact across DOE:

- FES: MFE (disruptions, plasma turbulence), magnetic reconnection
- NNSA: ICF (kinetic effects in hohlraum)

SNL-LANL productivity >130 papers since 2006

Staffing: ~1 FTE (Chacon, Burby, Chen, Koshkarov,Stanier)

Status: to be renewed in FY23

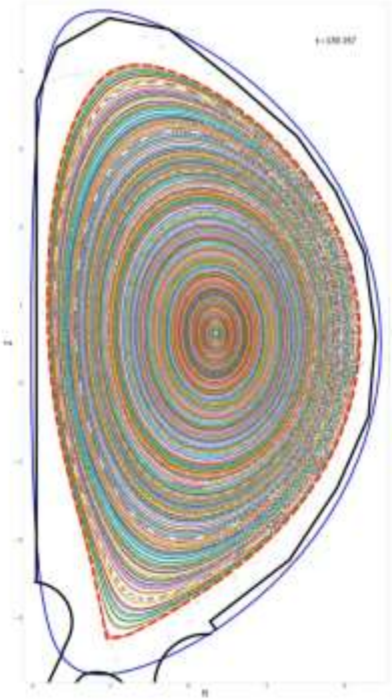

Evolution of magnetic field in a 3D ITER tokamak simulation of a (1,1) kink mode with realistic heat-transport anisotropy, $\frac{\chi_\parallel}{\chi_\perp} = 10^7$. As the (1,1) mode grows the core splits and gets expelled while magnetic stochasticity sets in at the outer edge, while the core eventually reheals to form a transport barrier.

Enabling Predictive Plasma Simulations
Historical perspective, partners

- Some **history**:
  - Project born in 2006 as an SNL-LANL collaboration on solvers, numerical discretizations, and applications of fluid and kinetic plasma-based PDEs
  - Project continued through Chacon's move to ORNL in 2008 and back to LANL in 2012
  - Has enabled many key technical breakthroughs over the years (see previous slide)
- Has impacted **several DOE agencies**, including:
  - DOE-FES Basic Plasma research (FY13-15 and FY16-18) on fast magnetic reconnection
  - DOE-FES-ASCR SciDAC (FY18-22):
    - Tokamak Disruption Simulation (TDS) Center: fully implicit MHD solvers
    - High-fidelity Boundary Plasma Simulation (HBPS) Partnership (Chacon ASCR Project PI): fully implicit particle-based kinetic solvers
  - DOE-NNSA (LDRD DR, FY21-23): High-fidelity simulation of ICF/HED experiments (Chacon, PI)
- Has nurtured **long-standing national and international collaborations**:
  - National: ORNL, LLNL, PPPL
  - International: Consorzio RFX, Padova, Italy (Reversed Field Pinch experiment)

Building linear structures into neural ODE stabilizes long-time predictions (PI: Q. Tang, $400K/yr)

Scientific Achievement

Developed a novel network architecture to stabilize neural ODEs, showing improved accuracy in long-time statistics.

Significance and Impact

The proposed network can produce much more accurate and robust surrogates for climate and weather predictability and accelerator dynamics. Both applications are being studied by our team through collaborations with domain scientists.

Research Details

- We propose to separate linear and nonlinear portions in vanilla neural ODE and build physics-motived structures, e.g., a linear convolution neural network (CNN), into the linear part of the network
- We demonstrate long-time stability is improved using PDE examples and show significantly improved accuracy in long-time statistics
- A novel ROM framework based on eigen-decomposition of learnt operators is proposed

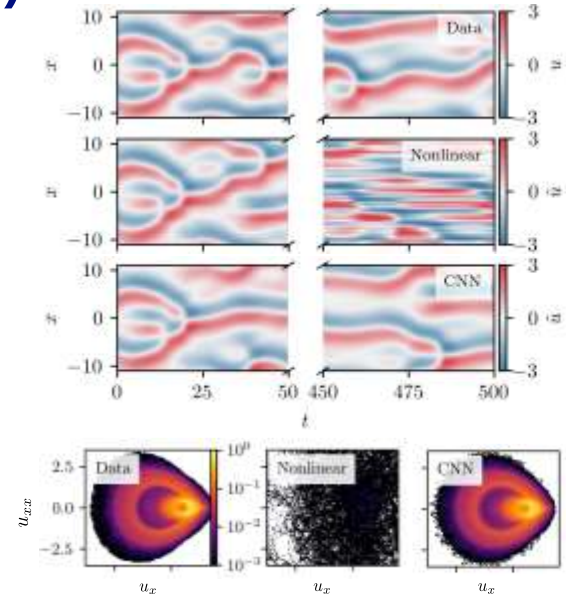

Top: predictions using different neural network architectures show our networks (CNN) predict significantly better than vanilla neural ODE (nonlinear), which becomes unstable at long time **Bottom:** a joint PDF of the first and second derivative show trajectories predicted by our network stay on the attractor at long time

A. Linot, J. Burby, Q. Tang, P. Balaprakash, M. Graham, and R. Maulik, submitted, arXiv:2203.15706, 2022

LDRD Investments in AM areas ~$22M/year in new money

Machine Learning: ~15-20 projects with $15M/yr in new money (crosscut topic)

Multiscale methods ($4.5M/yr)

High-fidelity Electromagnetic Simulation Capability for ICF/HED Experiments	Luis Chacon	20210063DR
Multiscale Method for Fluid-Kinetic Coupling in Ocean-World Jet Simulations	Daniel Livescu	20210298ER
MASS-APP: Multi-physics Adaptive Scalable Simulator for Applications in Plasma Physics	Daniel Livescu	20220104DR
High-Order IMEX Time Integration for Radiation Transport and Coupled Multiphysics Problems	HyeongKae Park	20220174ER
Kinetic Electron and Radiation Transport in Inertial Confinement Fusion-capsule Implosions	Andrei Simakov	20220200ER

Discretization ($1.5M/yr)

Adaptive high-order finite element ALE methods for multi-material hydrodynamics	Jacob Waltz	20200201ER
Mimetic Tensor-Train Algorithms for High-Dimensional PDEs without the Curse of Dimensionality	Gianmarco Manzini	20210485ER
Stochastic Finite Volume Method for Robust Optimization of Nonlinear Flows	Svetlana Tokareva	20220121ER
Predictive Computational Framework for the Treatment of Dynamic Fracture Problems on Polytopal Meshes	Hasem Mourad	20220129ER

Optimization (not ML; $0.7M/yr)

Accelerating Combinatorial Optimization with Noisy Analog Hardware	Carleton Coffrin	20210114ER
Fast, Linear Programming-Based Algorithms with Solution Quality Guarantees for Nonlinear Optimal Control Problems	Kaarhik Sundar	20220006ER

Spectral Plasma Solver (SPS): A multiscale framework for plasma simulation

G. L. Delzanno, O. Koshkarov, D. Svyatsky (T-5)

- **The problem:** Huge scale separation in time and space + high dimensionality (6D+time) make large-scale microscopic/kinetic simulations of many plasma systems simply impossible
- **The goal:** large-scale/fluid simulations that seamlessly incorporate microscopic physics only where necessary
- **The approach: Hermite spectral expansion methods**
  - SPS is based on a spectral approach: low-order terms of the expansion give the macroscopic/fluid behavior of the plasma while microscopic physics is retained by adding more terms to the expansion
  - It features flexible spatial and temporal discretizations: (i) spectral Fourier method or arbitrary order discontinuous Galerkin method; (ii) fast adaptive explicit or implicit time integration
  - SPS has already been successfully applied to problems where traditional methods struggle. Example: first numerical confirmation of a new kinetic turbulence regime near the Sun, called inertial kinetic-Alfven wave turbulence (Fig. 2)
  - To arrive at a flexible, adaptive, multi-scale framework for plasma physics we need: - The spatial and temporal adaptivity of the spectral expansion - Optimization of the spectral expansion - Exascale-readiness

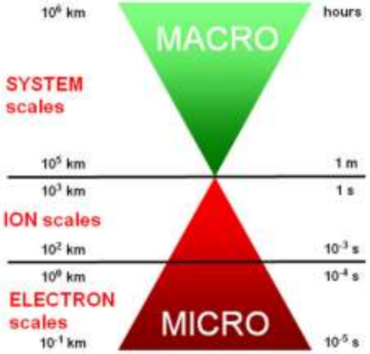

Fig. 1: scale separation in the tail of the Earth's magnetosphere

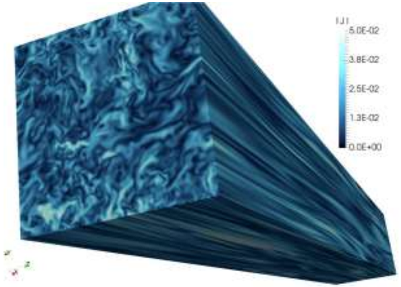

Fig. 2: 3D3V SPS simulation of inertial kinetic-Alfven turbulence

SPS historical perspective and connections

- FY13-15 ($300k/year, LDRD): SPS development started (fully spectral formulation with fixed number of spectral terms).
- FY17-19 ($300k/year, LDRD): Further development of SPS with focus on optimization and adaptivity.
- FY20-22 ($2.1M total, NASA): Beam-PIE proposal awarded. Objective: rocket experiment to demonstrate plasma waves generated by electron beam in space. SPS was critical for prediction of new highly efficient beam-plasma interaction regime motivating the experiment.
- FY22-25 ($1.7M/year, LDRD): to develop the next-generation, exascale-ready methods for multiscale plasma physics modeling. Based on SPS.
- FY22-25 ($1M/year, NASA): NASA strategic capability proposal selected. Objective: micro-macro modeling of the Earth magnetosphere based on SPS.
- Other projects in space physics portfolio enabled by SPS: 1) Modeling and optimization of radiation belt remediation; 2) Investigation of plasmapause stability; 3) Accurate modeling of wave-particle interaction physics.

Advanced Network Science Initiative (ANSI): Advancing the State-of-Practice in Energy System Modeling Through Fundamental Applied Mathematics

PI: Russell Bent (https://lanl-ansi.github.io/)

Complex, networked energy infrastructure systems provide a rich source of challenging problems that are ideally suited for applied mathematics solutions. ANSI develops novel techniques in optimization, control theory, machine learning to support DOE office of electricity (OE) programs like Advanced Grid Modeling (AGM), North American Energy Resilience Model (NAERM), Grid Modernization Laboratory Consortium (GMLC), and Microgrids R&D.

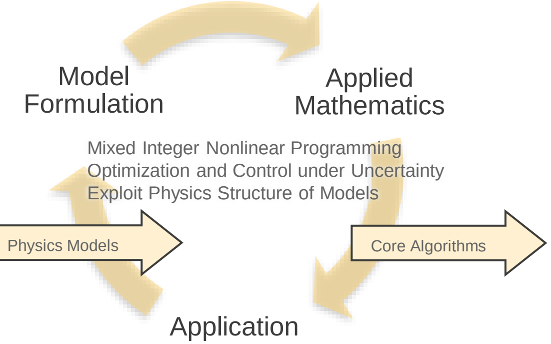

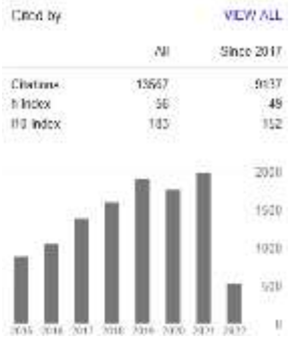

Team Citations. Source: Google Scholar

Key historical events in LANL's infrastructure science R&D. Highlights application (DHS) to fundamental theory (LDRD) to application (DOE/DHS/DTRA) scientific cycle.

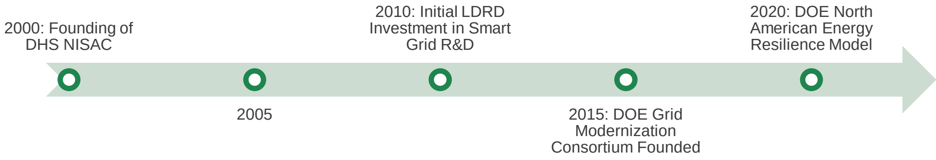

Research Approach: Identify the key applied mathematics gaps in computationally hard-to-solve energy problems, develop a fundamental algorithm to solve the problem, extract and generalize the algorithm's key features for the scientific community and to directly impact the application domain.

Highlight: New global optimization algorithm that a) **outperformed** literature by up two orders of magnitude on academic benchmarks and b) energy problems https://github.com/lanl-ansi/Alpine.jl

Highlight: New method for addressing convergence challenges in solving energy infrastructure nonlinear, algebraic equations. **Multiple Contingency Solver** succeeds > 99% of the time. https://github.com/lanl-ansi/InfrastructureModels.jl

Highlight: Organize bi-annual Grid Science Winter School and Conference to train the next generation of applied mathematicians for infrastructure R&D.

AI Institute for Advances in Optimization

Collaborative Network

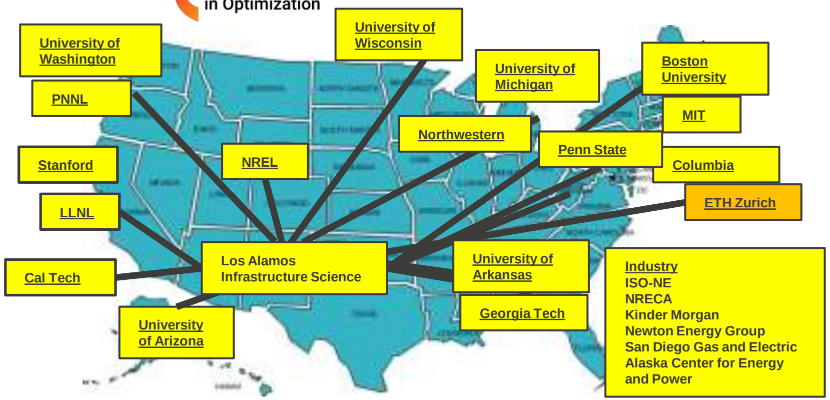

LDRD DR: The Optimization of Machine Learning: Imposing Requirements on Artificial Intelligence for High Consequence Decision Making PI: Russell Bent

Machine learning (ML) methods often lack the ability to directly incorporate requirements and **meaningful theoretical guarantees that are necessary for high profile, high-consequence decision problems**, such as the physics-informed science and security applications entrusted to the DOE. This project develops **theoretical and algorithmic advances** in mathematical optimization for **constrained learning (CL), robust learning (ROL), and tractable learning (TL) to meet these requirements**.

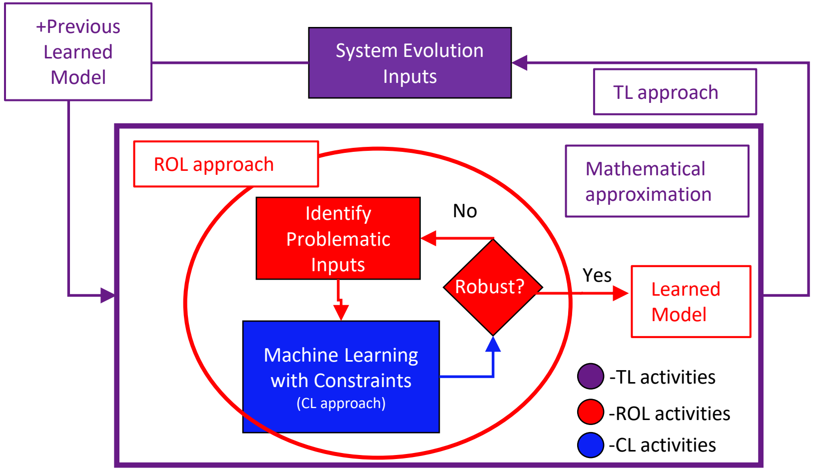

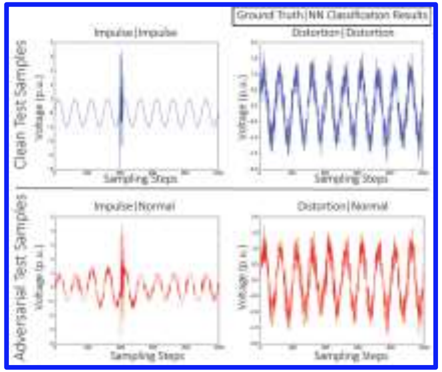

Example showing failure of traditional ML strategies (and success of TL) to identify abnormal behavior in the presence of noise (red, abnormal; blue, normal)

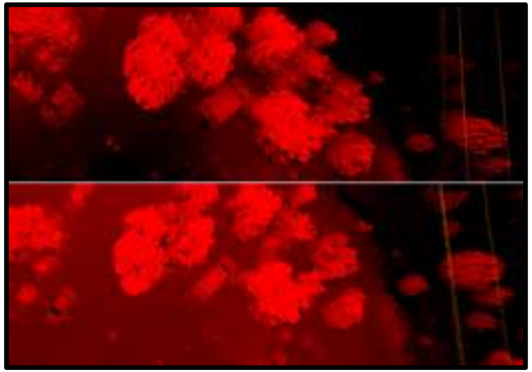

Example of TL in action in determination of unmanned vehicle roll, pitch, and raw when taking LIDAR measurements (top, raw; bottom, TL-cleaned)

Productivity: Papers in NeurIPS, ICML, Journal of Global Optimization, and other ML and optimization venues. Applications in energy systems and autonomous vehicles with papers in the relevant engineering domains. (24 papers in 2020 and 2021)

Transformational Research on High-Order Methods for Unstructured Polytopal Meshes

M. Shashkov (XCP-4), K. Lipnikov (T-5), E. Kikinzon (CCS-7), S. Tokareva (T-5)

High-order interface reconstruction methods (1 MS in-prep)

MOF[2]: We are developing the **Moment-Of-Fluid** method based on the 2nd-order moments. It is superior to existing state-of-the-art methods in capturing accurately fine-scale features of different materials in a single cell.

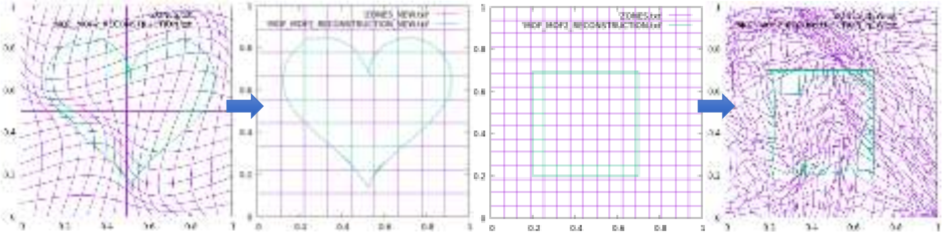

Remapping of interfaces between arbitrary meshes

High-order methods for multi-material cells

New methods (developed partially under ASCR program) for solving elliptic PDEs on polytopal meshes non-aligned with material interfaces replace standard homogenization approach and account for subscale interactions.

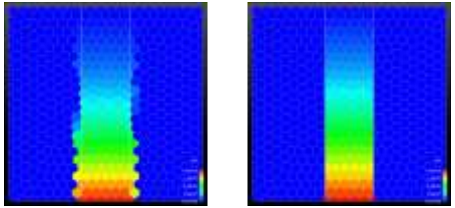

Homogenization ASC

Heat propagation in the channel with a slow cross diffusion: incorrect behavior (left) is fixed with the new method (right).

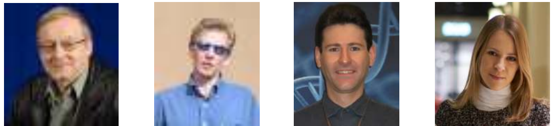

High-order hydrodynamic methods (3 papers since 2017)

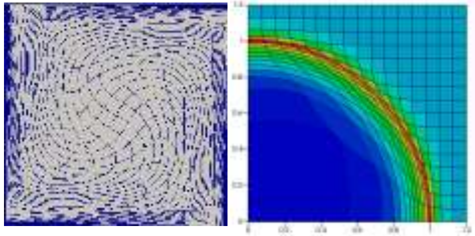

Taylor-Green vortex, T=1.25 Sedov explosion, T=1

New **compact** high-order **multidimensional upwind methods** for advection-diffusion PDEs that have superior accuracy and efficiency on emerging computing architectures.

High-order conservative data transfer methods (1 paper)

Algorithms for conservative bounds-preserving remap of momentum between unstructured meshes are used in new ASC Lagrangian codes.

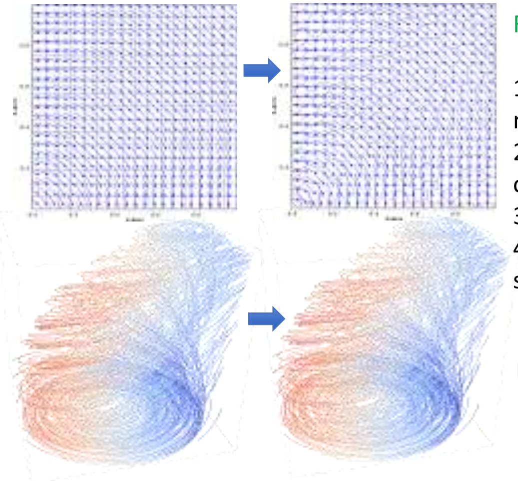

Remap for Cell-Centered Hydro:

1. Compute density and specific momentum in source mesh cells
2. Computer linear reconstructions on the source mesh
3. Intersect meshes
4. Integrate over target mesh subcells

Error is about or less than 0.2%

LDRD DR + ASC: Initiative for Scientific Imaging (ISI)

C. Garcia-Cardona (CCS), M. Klasky (T), Y. Lin (EES), S. Vogel (MST), B. Wohlberg (T)

- **Goal**: to solve difficult imaging inverse problems central to LANL's (and DOE's) mission.
- Applications to **materials science** (numerous imaging modalities), **stockpile stewardship** (DARHT, pRAD), **non-proliferation and energy security** (seismic imaging).
- **Current funding:** LDRD DR (Prioritizing the Prior, PI: Wohlberg), NNSA/NA-80/ASC (PI: Klasky)
- **Research foci**:
  - Development of advanced priors based on proximal optimization algorithms (cf. award-winning Plug-and-Play Priors framework co-developed at LANL) and deep learning for scientific imaging problems
  - Deep unrolling of optimization algorithms and stochastic gradient methods for large-scale problems and near-real-time solutions
  - Inversion algorithms that are robust to errors and uncertainties in the forward model
- **Achievements/Impact**:
  - New state-of-the-art algorithm for ptychographic imaging (national security and materials science applications)
  - Significantly extended neutron imaging capability at LANSCE. Contributed to first characterization of irradiated nuclear fuel (see Fig.)
  - Significant progress in developing physics-informed machine learning approach for hydrodynamics experiments (DARHT)
  - Advanced machine-learning methods for seismic imaging
  - Open-source code designed to make advanced techniques accessible to LANL community released Q4 2021

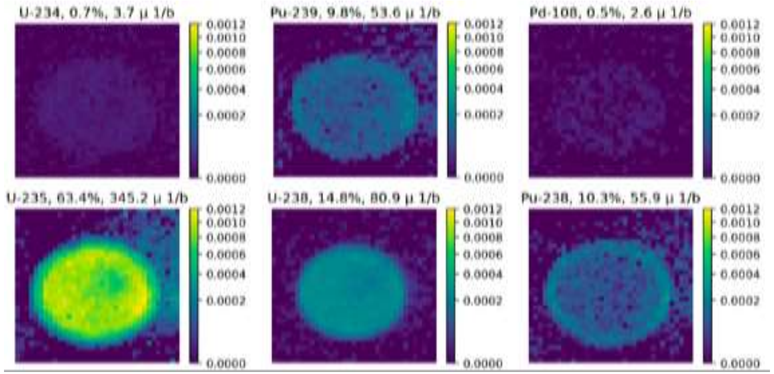

Energy-resolved neutron imaging of a 6mm diameter irradiated U-10Zr-1Pd sample providing spatial distributions of various isotopes.

Los Alamos NATIONAL LABORATORY

Computational Partnerships

Galen Shipman, Applied Computer Science,
Computer, Computational, and Statistical Science Division

May 12, 2022

Computational Partnerships - SciDAC

PI	Lead	Title	Offices	Total funding amount	Collaborators
Jones, Phil	LANL	Coupling Approaches for Next Generation Architectures (CANGA)	BER/ASCR	$3.5M / $2.0M	SNL, ANL, ORNL, UC-Davis, Stony Brook, Ohio State
Price, Stephen	LANL	Probabilistic Sea-Level Projections from Ice Sheet and Earth System Models	BER/ASCR	$4.6M / $689K	LBNL, ORNL, SNL, NYU, U Mich, UC Irvine
Turner, Adrian	LANL	A New Discrete Element Sea-ice Model for Earth System Modeling	BER/ASCR	$1.3M / $0	SNL, Naval Postgraduate School
Tang, Xianzhu	LANL	Tokamak Disruption Simulation	FES/ASCR	$4.0M / $750K	PPPL, LLNL, SNL, Virginia Tech, SNL, Columbia Univ, ANL, U of Texas-Austin, Univ Maryland
Andersson, Andres (David)	**LANL**	**Simulation of Fission Gas in Uranium Oxide Nuclear Fuel**	**NE/ASCR**	**$1.6M / $0**	**ANL, INL, ORNL, SNL, U Florida, U Tenn-Knoxville**
Carlson, Joseph & Lawrence Earl	LANL	Nuclear Computational Low Energy Initiative (NUCLEI)	NP/ASCR	$887K / $550K	LBNL, RNL, Iowa State, Indiana U, LLNL, Mich State, NCSU, ORNL, OSU, UNC, Univ Oregon, UTK
Perez, Danny	**LANL**	**FusMatML: Machine Learning Atomistic Modeling for Fusion Materials**	**FES/ASCR**	**$900K / $0**	**SNL, UTK**
Scheinker, Alexander	SLAC	Leveraging ML/AI techniques to enable a breakthrough in ultra-short-bunch	HEP/ASCR	$0/$40k	SLAC
Moulton, John David	LBNL	Exasheds: Advancing Watershed System Understanding through Exascale Simulation and Machine Learning	BER/ASCR	$0/$315k	LBNL, PNNL, ORNL, USGS, UT Austin
Ahrens, James	**ANL**	**RAPIDS2: A SciDAC Institute for Computer Science, Data, and Artificial Intelligence**	**ASCR**	**$927K**	**ANL, BNL, LBL, LLNL, ORNL, Northwestern, University of Deleware, The Ohio State University, University of Oregon, Rutgers, University of Florida, Kitware**
Karra, Satish	LBNL	Development of Terrestrial Dynamical Cores for the ACME to Simulate Water Cycle	BER/ASCR	$294K / $0	LBNL
Nadiga, Balu	SNL	Non-Hydrostatic Dynamics with Multi-Moment Characteristic Discontinuous Galerkin Methods	BER/ASCR	$875 / $0	SNL
Chacon, Luis	PPPL	High-Fidelity Boundary Plasma Simulation (HBPS): ML strategies for particle collisions in plasmas (Chacon)	FES/ASCR	$0 / $1.0M	ORNL, CU-Boulder, LBNL, MIT, RPI, Rutgers, UT-Austin, UC-San Diego, SNL
Martinez-Saez, Enrique	U Tenn	Plasma Surface Interactions: Predicting Performance and Impact of Evolving PFC Surfaces	FES/ASCR	$1.5M / $450K	U Tenn-Knoxville, LLNL, RPI, UC-San Diego, U Mass, PNNL, U Illinois, U Missouri
Lawrence, Earl	**ANL**	**Accelerating HEP Science: Inference and Machine Learning at Extreme Scales**	**HEP/ASCR**	**$0 / $500K**	**ANL, BNL, LBNL, Virginia Tech**
Fryer, Christopher	ORNL	Towards Exascale Astrophysics of Mergers and Supernovae (TEAMS)	NP/ASCR	$540K / $0	ORNL, Stony Brook, UC-Berkeley, Mich State, LBNL, UC-San Diego, Univ Tenn, Notre Dame, Princeton, ANL, U WA
Yoon, Boram	TJNAL	Computing the Properties of Matter with Leadership Computing Resources	NP/ASCR	$0 / $300K	GWU, Mich State, PNNL, U WA, BNL, Intel, MIT, NVIDIA, UNC, William & Mary

Simulation of Fission Gas in Uranium Oxide Nuclear Fuel (Andersson)

- **Project goal:** Advance the understanding of fission gas behavior in UO_2 by developing a meso-scale fission gas simulator accounting for atomic scale mechanisms that takes advantage of leadership class computers and UQ approaches.
- **LANL role:** Project leadership, identifying mechanisms of gas transport, focused on new atomistic methods based on DFT/MD/AMD to characterize defect mobilities and cluster dynamics for upscaling to higher level models, and application of UQ to the proposed transport mechanisms.

Atomic scale mechanisms

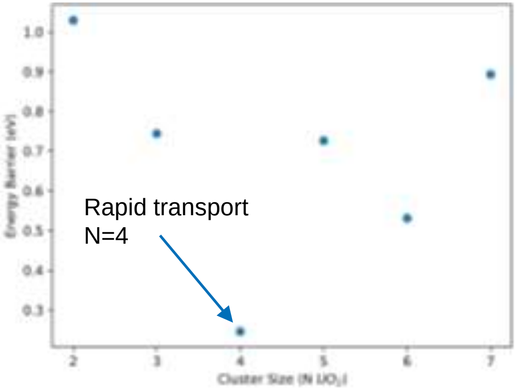

MD simulations of diffusion of interstitials ($N \times U_i 2O_i$). Arrhenius barrier has a a very low minimum at N=4.

Uncertainty quantification

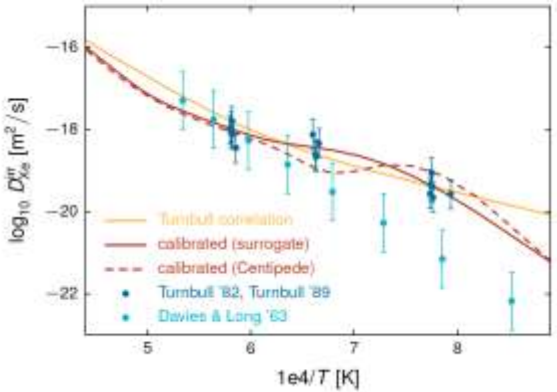

Maximum a posteriori (MAP) Xe diffusivity obtained by Bayesian calibration of the 183-parameter Centipede cluster dynamics code.

Meso-scale simulation

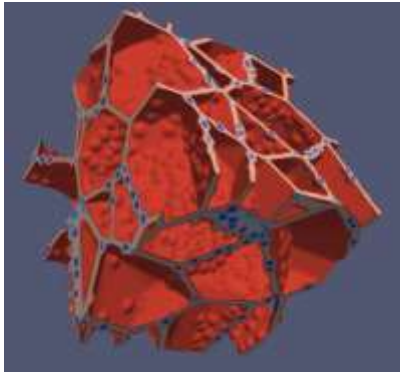

Coupled Xolotl-Marmot meso-scale simulation of gas evolution in 3D resolving 20 grains.

FusMatML: Machine Learning Atomistic Modeling for Fusion Materials (Danny Perez, Nicholas Lubbers)

Development of Machine Learning methods for fusion materials

Recently completed:

- Demonstrated automated generation of W potential using information entropy optimization
  arXiv:2201.09829 (submitted)

- Developed MLIAP-Unified: links LAMMPS with python-enabled ML ecosystem based on pytorch, tensorflow, and jax

Work in Progress:

- Correlated error model for ML potentials
- Chemical-embedding SNAP formalism for multi-component systems

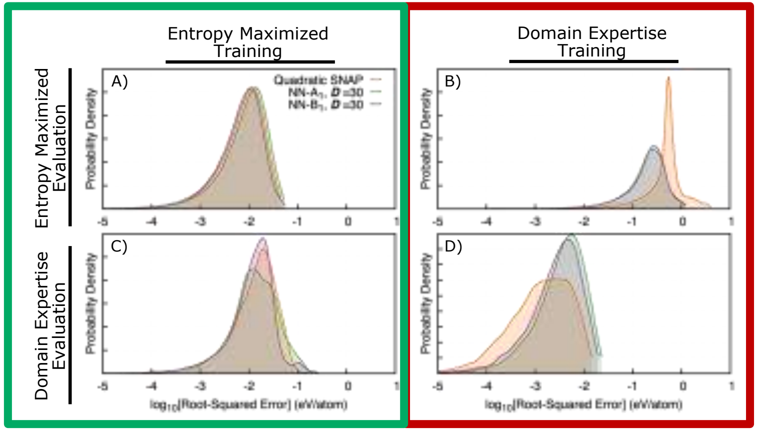

Improved method: slight loss in local accuracy, but extremely robust

Conventional method: locally accurate but not transferable

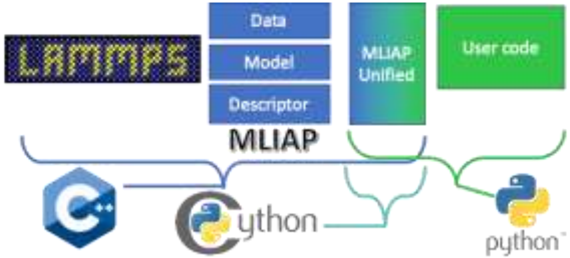

LANL's Lightweight In Situ Visualization & Production Visualization Activities in SciDAC RAPIDS

Scientific Achievement - created Seer-Dash which enables in situ visualization and analysis while requiring only very lightweight software (Mochi & python).

Significance and Impact - Resource allocation on HPCs and heavy requirements of software like ParaView & VisIt is often a barrier to in situ analysis. Seer-Dash lowers this barrier by providing a simple connection interface from the supercomputer to the user's laptop for analysis.

Research Details

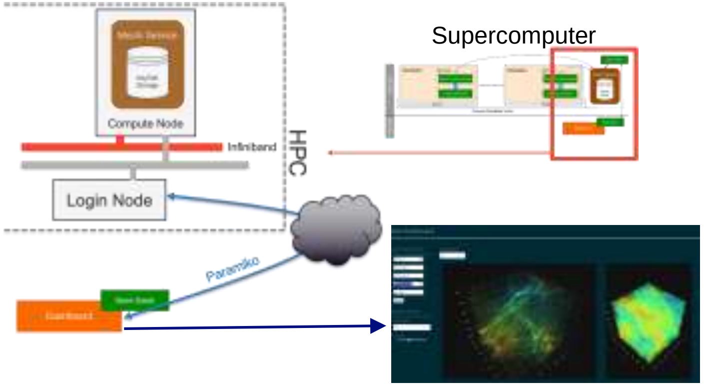

Sea Ice/Level Visualization

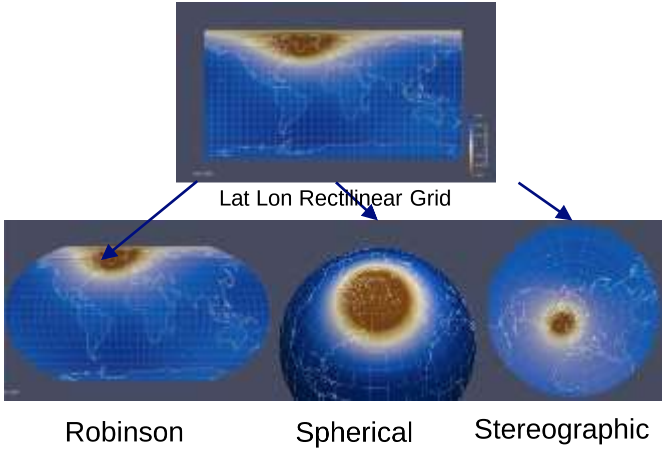

Supported BER scientists studying Sea Level Change by developing and delivering filters that allow them to change from Lat/Lon coordinates to a Robinson, spherical, and stereographic projection interactively while using ParaView.

Pascal Grosset and James Ahrens. 2021. Lightweight Interface for In Situ Analysis and Visualization of Particle Data. *ISAV'21: In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization*. https://doi.org/10.1145/3490138.3490143

PI: James Ahrens

Los Alamos

Accelerating HEP Science: Inference and Machine Learning at Extreme Scales (Lawrence)

- **Scientific Achievement**: Updates to the CosmicEmu, which produces highly accurate predictions of matter power spectra without the need for additional intensive simulations.
  - Built with SEPIA, a new Python tool for emulating and calibrating large-scale simulations
  - Corrected data pipeline issues (mis-labeling of red shift, alignment of simulation resolutions)
  - Parallelized main function to enable faster emulation of multiple cosmologies / red shifts
- **Significance and Impact**: The CosmicEmu can give predictions of the nonlinear matter power spectrum in fractions of a second, instead of the days required to produce the same result from simulation. This can be used to accelerate data analysis pipelines for surveys like those at the upcoming Vera C. Rubin Observatory.

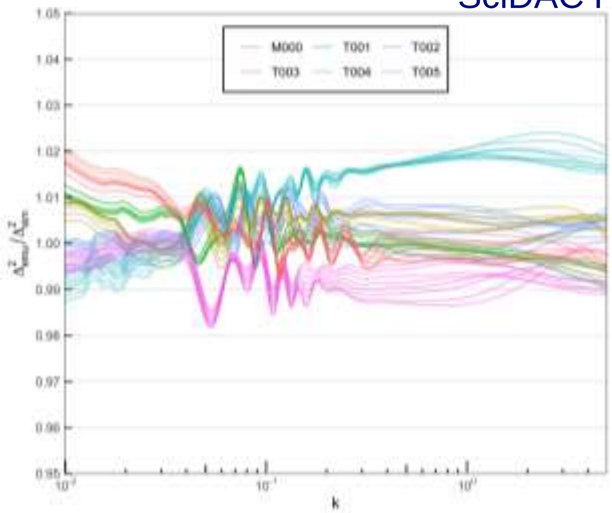

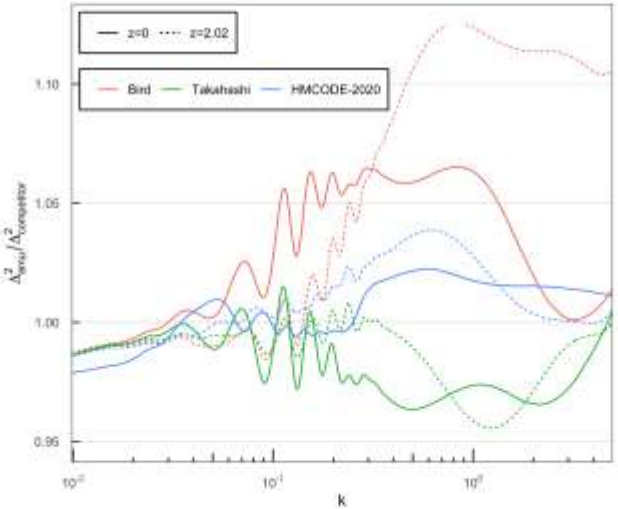

Top: Prediction results from the trained emulator for a test set of simulation results. The emulator has about 2% error.
Bottom: Comparison of emulator output to commonly used competitors. There is relatively higher agreement with HMCODE-2020.

Multi-Physics: North American Energy Resilience Model (NAERM) Russell Bent

Multi-lab Partnership (LLNL, PNNL, ORNL, LANL, SNL, INL, ANL, NREL)

Vision - Rapidly predict energy system interdependencies, consequences and responses to extreme events at a national scale

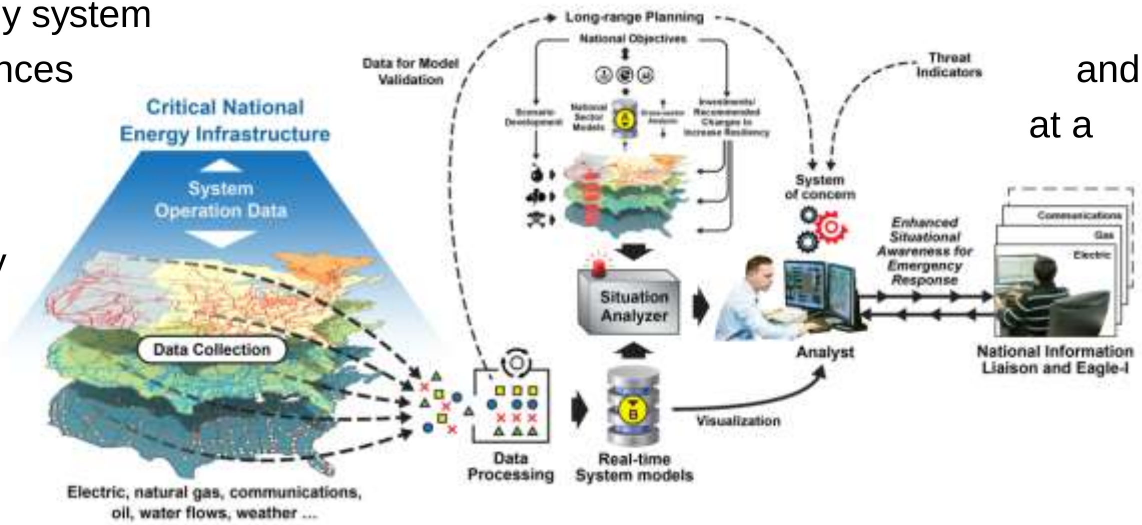

Mission - Develop and deploy engineering-class modeling system for planning and real-time resilience analysis

Key Objective – Catalyze partnerships with industry, national labs, states/communities and other federal agencies to enhance coordination to support energy resilience

North American Energy Resilience Model

LANL Contributions

Strategic and Technical Leadership

Russell Bent, T-5

- Contributed to DOE vision for NAERM
- Area lead for pipeline modeling efforts (LANL, ANL) and validation, verification, and uncertainty quantification efforts (LANL, SNL, ORNL)

Pressure and flow profiles that track system evolution and impacts over time

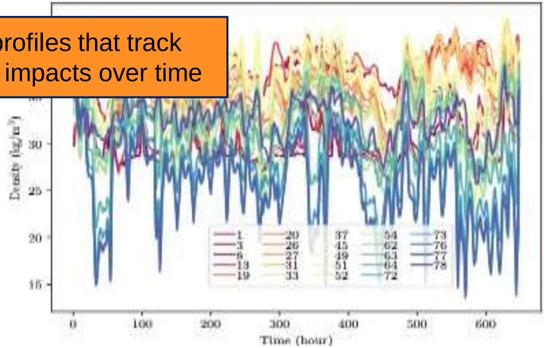

Physics Modeling and Applied Mathematics

- GasModels.jl
  - Open source software for steady state modeling and optimization of compressible fluids for long term planning
- Grail (Gas Reliability Analysis Integrated Library)
  - Open source software for transient modeling of control, market modeling, and simulation of compressible fluids in pipes
- Integration of multi model physics
  - Gas, electric, communication networks

Map depicts interconnected models for the Kern River and Questar pipelines

Los Alamos NATIONAL LABORATORY

Computer Science R&D

Pat McCormick, Applied Computer Science,
Computer, Computational, and Statistical Science Division

May 11, 2021

Legion Programming System Update

- ECP Efforts:
  - Support in place for AMD GPUs; S3D "Regent" running
  - Focused on scaling >= 6,000 ranks
  - Outreach from Intel – Intel-wide invited talk in April '22
- DISTAL: DSL for dense tensor algebra on distributed and heterogeneous systems.
  - Independently describe how tensors and computation map onto target machines
  - Out- performs existing systems by between 1.8x to 3.7x (with a 45.7x outlier) on higher order tensor operations.
  - Wrapping up ASCR effort w/ RL-based mapping
- Continued engagement with external community
  - **NVIDIA:** "cuNumeric" and other products
  - At the cost of some of our staff @ LANL…
  - Growing interest from other customers

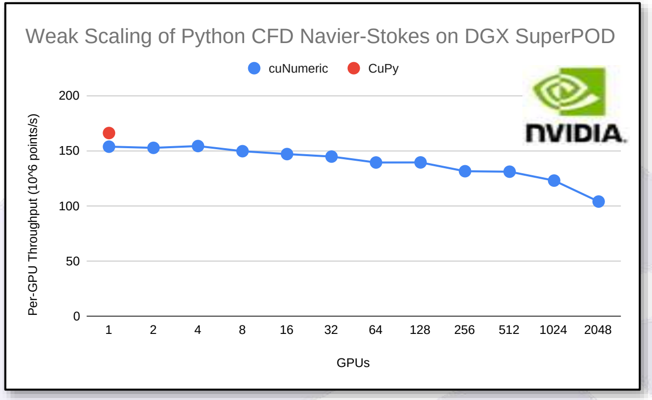

Publications

- Index Launches: Scalable, Flexible Representation of Parallel Task Groups: R. Soi, M. Bauer (NVIDIA), S. Treichler (NVIDIA), M. Papadakis (NVIDIA), W. Lee (NVIDIA), P. McCormick (LANL), A. Aiken (Stanford), E. Slaughter (SLAC), The International Conference for High Performance Computing, Networking, Storage and Analysis (SC21), Nov. 14-19, 2021, St. Louis, MO.

- DISTAL: The Distributed Tensor Algebra Compiler, R. Yadav, A. Aiken, and F. Kjolstad, Proceedings of the Conference on Programming Language Design and Implementation (PLDI '22), June 13-17, 2022, San Diego, CA.




SLAC NATIONAL ACCELERATOR LABORATORY NVIDIA. ECP EXASCALE COMPUTING PROJECT PSAAP III

FlexFlow: Joint Optimization of Algebraic Transformations and Parallelization

- Part of our Legion ECP efforts
  - Today’s ML frameworks are limited by approaches to parallelism (e.g., static/fixes & tailored/custom)
  - New support for joint graph and parallelism optimizations (OSDI ‘22)
  - AMD GPU support nearing completion
  - Long term: mixed precision support

Traditional Framework Design

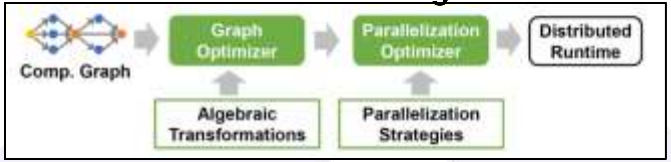

FlexFlow “Unity” Design

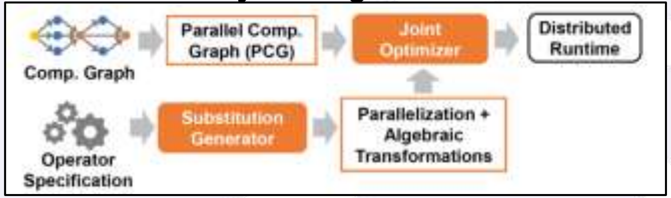

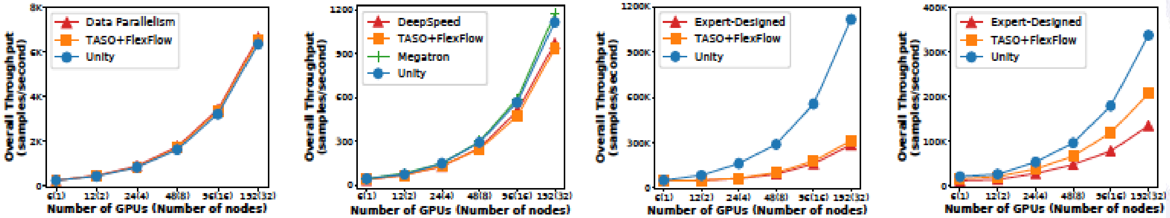

(a) ResNeXt-50. (b) BERT-Large. (c) DLRM. (d) CANDLE-Uno.

- Staffing woes are unfortunately a growing concern. Legion-related efforts have lost 3 key staff to industry…

Publications:

- Unity: Unity: Accelerating DNN Training Through Joint Optimization of Algebraic Transformations and Parallelization: Colin Unger (Stanford), Zhihao Jia (CMU), Wei Wu (LANL/NVIDIA), Sina Lin (Microsoft), Mandeep Baines (Facebook), Carlos Efrain Quintero Narvaez (Facebook), Vinay Ramakrishnaiah (LANL), Nirmal Prajapati (LANNL) Patrick McCormick (LANL), Jamaludin Mohd-Yusof (LANL), Xi Luo (SLAC), Dheevatsa Mudigere (Facebook), Jongsoo Park (Facebook), Misha Smelyanskiy (Facebook), Alex Aiken (Stanford), 16th USENIX Symposium on Operating Systems Design and Implementation, July 11–13, 2022, Carlsbad, CA.

Compiler Scope: LLVM Foundations

- Initial LLVM efforts supported by ASCR R&D and CSGF (James Jablin) circa 2011.
- NNSA ECP: “Kitsune+Tapir” w/ MIT (Shardl & Leiserson)
  - Rethink compiler design for parallelism
  - Generalize parallel constructs, not as a front-end, language/programming system specific design
  - Beating hand-coded CUDA and Kokkos by +2x, OpenMP offload by +4x, seeing 5-17% boost on CPU targets
  - Expanding feature set for more forms of parallelism and architectures
  - Via NNSA PSAAP III – exploring use cases w/ Julia as well as advanced tooling for memory hierarchies.

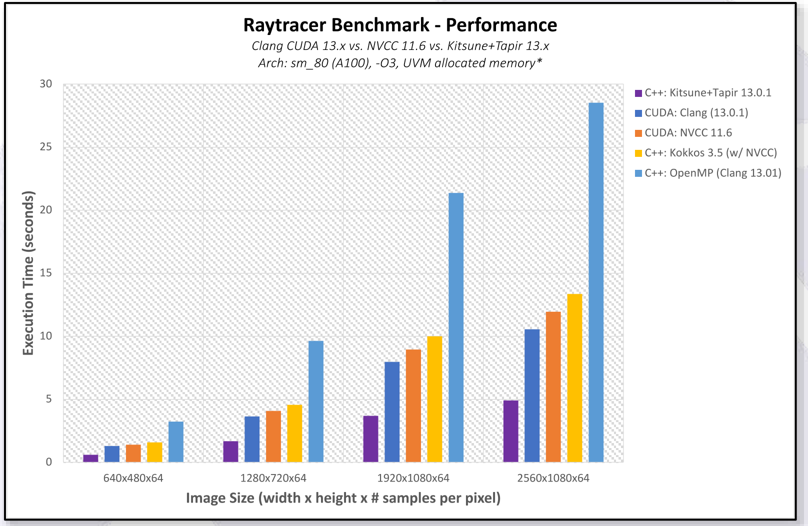

Note: Kokkos and OpenMP numbers are based on explicit data movement calls/directives. Kitsune+Tapir figures are using UVM.





Flang: Community Supported Fortran for LLVM

- Initial NNSA investment, late to ECP:
  - LLVM-based Fortran front end
  - Vendors have product plans based on Flang
  - Community continues to (slowly) grow – however, more “bottlenecks” out of the way…
- Still lots of work to do:
  - Fortran 95 complete, finish through Fortran 2018
  - Testing, testing, testing…
  - Performance
- DOE staffing is likely largest risk at this point
  - We’ve lost staff across 3 of the 4 labs involved to industry
  - Production-quality compilers are a lot of work

NVIDIA arm Google Los Alamos NATIONAL LABORATORY Argonne NATIONAL LABORATORY OAK RIDGE National Laboratory BERKELEY LAB AMD Linaro HUAWEI

The LLVM Compiler Infrastructure

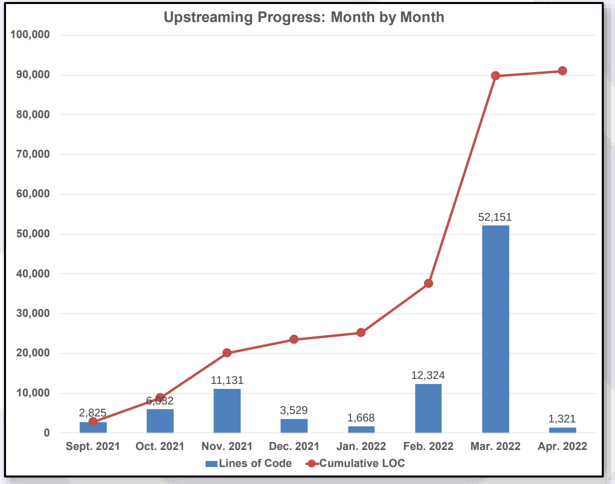



ECP EXASCALE COMPUTING PROJECT

Networking, Workflows and Edge Computing

Pat McCormick, Applied Computer Science,
Computer, Computational, and Statistical Science Division

May 12, 2022

5G Drones: Real Time Data Assimilation to Transform Wildfire Predictability

- **Research team:** Jon Reisner (PI), Dubey Manvendra (data discovery), Alex Josephson (fire modeling & partners), and Humberto Godinez (surrogate modeling)
- **Objectives:**
  - Use 5G drone technology to collect smoke/soot concentrations to inform and refine smoke production models within HIGRAD-FIRETEC
  - Examine utility of 5G drone data to predict fire spread and energy release rate
  - Interact and provide guidance with various partners utilizing drones in an operational setting
- **Expected Results:**
  - 5G drone data (visible & thermal images, smoke & soot information) will be invaluable with regard to identifying and calibrating chemical processes found within the smoke model
  - 5G drone data in combination with other data sources will lead to “gold” standard data sets that are invaluable for model validation
  - Knowledge transfer to partners will enable refined utilization of drones in an operational setting for enhanced predictions regarding fire spread & changes in intensity, e.g., utilization of surrogate model in an operational setting

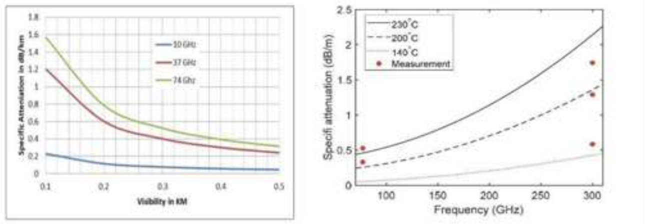

Fig 1a) Attenuation of 5G signals 10, 37 & 74 GHz as a function of smoke amount measured by visibility shows that the attenuation is largest at the shortest wavelength (Singh 2019) Fig 1b) Attenuation as a function of mm wave frequency modeled at 3 temperatures, reflecting more moisture in smoke at higher temperatures and bigger particles the red points are observations (Bystrov 21)

5G drones in combination with new modeling approaches will lead to better predictions of smoke and fire spread

Workflows for LANL Experimental Science

Projects: LDRDs, ExaFEL, ExaLearn, RAVEN (IARPA)
Scope: Accelerate user facility experiment workflows to provide higher quality data and thus shorter time to scientific discovery. Use data science, machine learning, and HPC to provide real-time results that can help steer experiments.
Staff: 10-12
Workflow Challenges:

- Increasing experimental data volumes and velocity
- Real-time analytics for experiment steering
- Workflow portability across facilities

Workflow Solutions:

- Machine learning for real-time inference
- HPC for real-time computation
- Statistical experimental design
- Fast data transfer
- Auto-updating visualization and data management

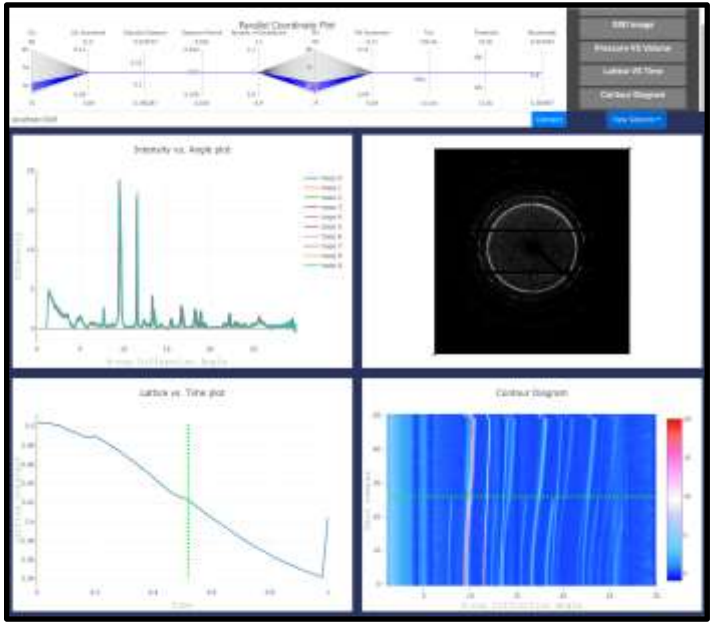

Real-time experimental workflow results displayed in Cinema:Snap visualization tool.

Example Workflows for LANL

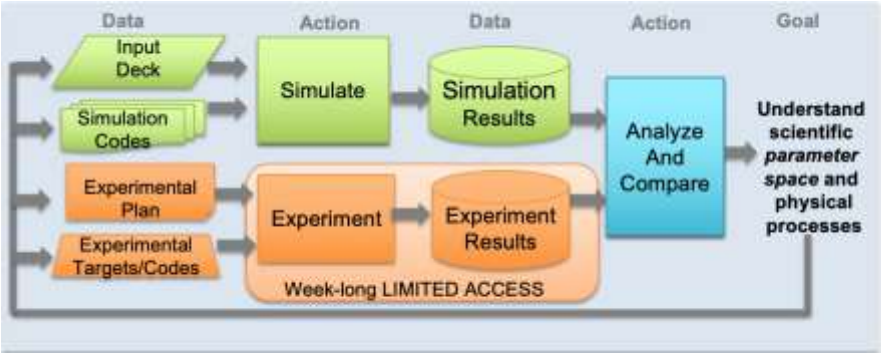

ASSIST LDRD: Integrated simulation/emulation for acceleration of real-time dynamic compression experimental workflows

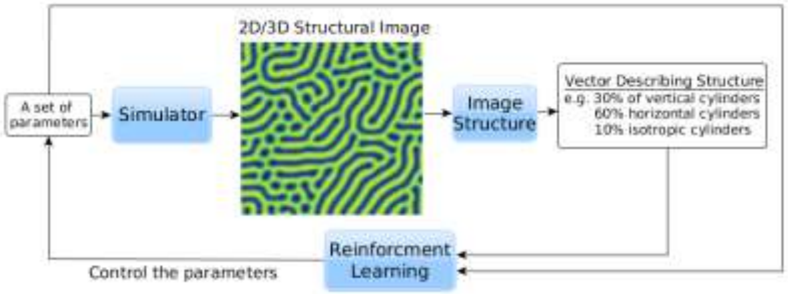

ExaLearn: Reinforcement learning workflows for control of self-annealing block copolymer experiments at light sources.

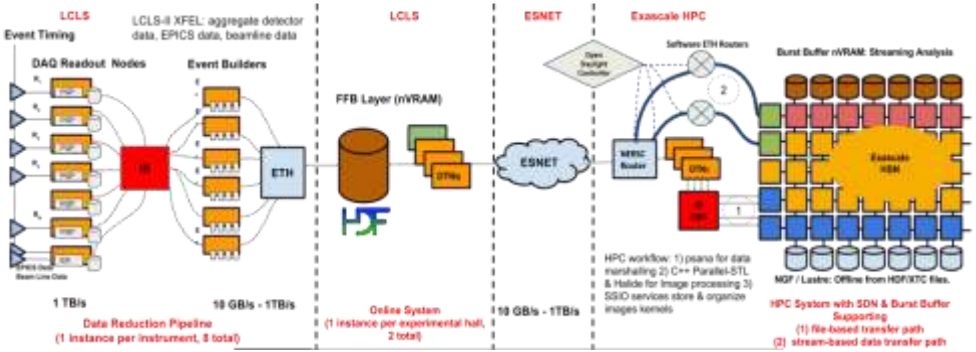

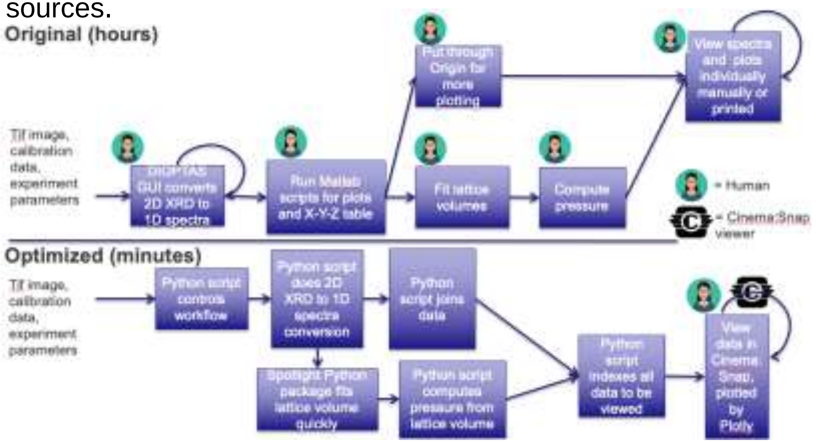

ExaFEL: Interfacility (between SLAC and NERSC) workflow for fast image reconstruction of crystallographic and single particle imaging data.

Mesoscale LDRD: Acceleration of rapid compression workflows using data science, machine learning and HPC along with real-time visualization.

Scientific Data

Pat McCormick, Applied Computer Science,
Computer, Computational, and Statistical Science Division

May 12, 2022

Scientific Data @ LANL

- Storage systems and I/O
  - **Mochi,** Lustre (2002), DeltaFS, HXHIM, Burst Buffer
  - Grand Unified File Index (GUFI) find files in less than a second from tens of billions of files at a site security as an administrator or user
  - Correlated Failure Consultation Tool for Operational Reliability (CoFaCTOR) two tier erasure protection against both random incident and correlated mass device failure
  - DeltaFS single dimensional indexing of output on the fly (Best Student Paper SC21)
  - Accelerated Box of Flash (ABOF) industry partnership producing the first computational storage standard based file system offload
  - Computational Storage industry partnerships providing hardware accelerated Ordered Key Value Store solutions for use in HPC parallel settings
- Data visualization and analysis
  - Paraview (1999, Dan Hitchcock), VTK, VTK-m, Catalyst
- In-Situ Inference
  - Advanced data science in exascale simulations
- Scientific Workflows
  - Large-scale ASC campaigns spanning multiple years running at scale (1/4 or more of Trinity)
  - Workflows for experimental science (ExaFEL, ExaLEARN, RAVEN – IARPA)
- Application engagement and community outreach
  - Delivering technologies through SciDAC (RAPIDS-2)

Best Student Paper SC 21
DeltaFS: A Scalable No-Ground-Truth Filesystem for Massively-Parallel Computing: Qing Zheng, Chuck Cranor, Greg Ganger, Garth Gibson, George Amvrosiadis, Brad Settlemyer, Gary Grider Carnegie Mellon University, LANL
Recent Patents
- Granted US Patent 11,080,196 Pattern-Aware Prefetching using parallel log-structured file system 8/3/2021
- Granted US Patent 11,048,699 Grand Unified File Indexing, 6/29/2021

.

Mochi: Rapid Development of Services for DOE Science

PI Rob Ross (ANL), co-PI Galen Shipman, LANL ASCR: $400K ($187K to CMU)

ASCR R&D (2015-2018)

- Methodology for rapid development of services for DOE science and platforms
- Demonstration of concepts in DeltaFS
- Early component prototypes

ASCR R&D, SciDAC, and ECP (2017-Present)

- Richer capabilities and AI opportunities (ASCR R&D)
- Use in HEP data analysis at ASCR facilities (SciDAC)
- Productization and software test (ECP)

Impact

- Leveraged by numerous DOE computer science teams
- Accelerating HEP experimental data analysis
- Enabling Intel DAOS, storage technology for Aurora system

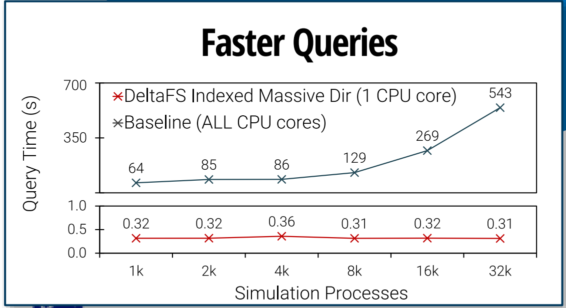

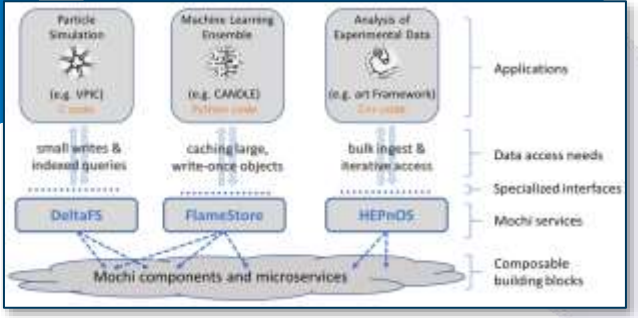

https://www.mcs.anl.gov/research/projects/mochi/

ASCR Mochi project accelerates data service development for DOE science and compute platforms

Scientific Achievement

Mochi has enabled numerous DOE computer science teams, and industry, to more rapidly build new data services through a thoughtful design methodology and reusable components.

Significance and Impact

Data services traditionally took many years to develop and productize. The Mochi project is shortening this development cycle, allowing teams to develop services specialized to their needs while still enabling significant component reuse.

Technical Approach

- Built using proven remote procedure call (RPC), remote direct memory access (RDMA), and user-level threading
- Define methodology for design of services using common components wherever possible
- Provide numerous typical capabilities via reusable components
- Exploring learning approaches to configuration and optimization

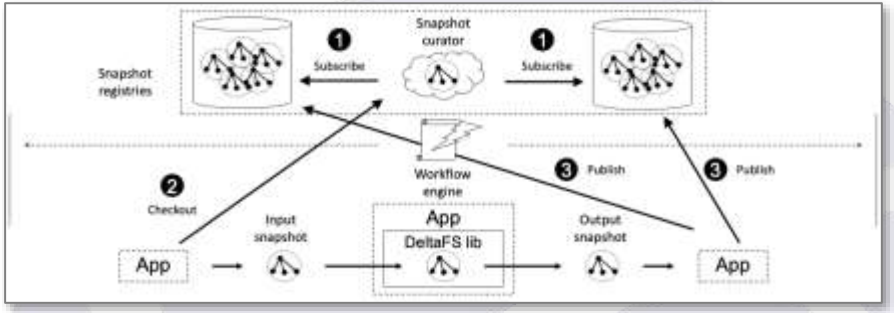

In science workflows, tasks (apps) often share data through a common service (e.g., a file system). In this work we are examining how to scale this type of service through more rich synchronization of this shared view.

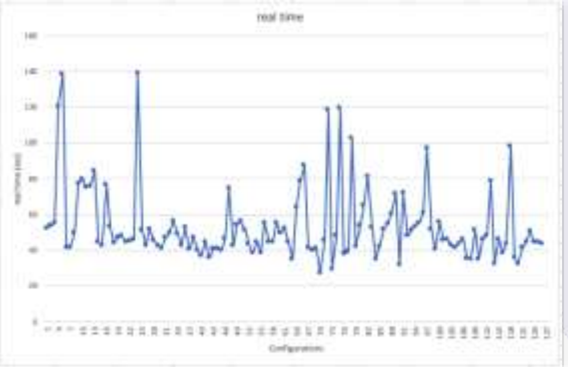

Given a workload, the DeepHyper tool efficiently explores the configuration parameter space for a Mochi service, building a surrogate model of this space to find high performing configurations. *With Jaehoon Koo, Sandeep Madireddy, and Prasanna Balaprakash (ANL)*

https://www.mcs.anl.gov/research/projects/mochi/

An entire ecosystem of vis technologies build upon long standing ASCR investments in LANL invented technologies

Cinema and workflow capabilities such as in situ feature detection and post processing algorithms, assist scientists in understanding data and making discoveries.

- Rogers (LANL) – PI and lead developer
- LANL staff: Biwer, Bujack, Dutta, Lo, Stam, Turton

`cinemasci` python module delivered & integrated into **ParaView** 5.9

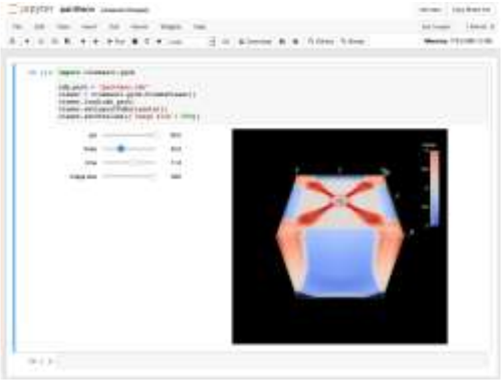

An ECP mini-app Cinema database can be viewed via a Jupyter notebook workflow

Integration with ECP infrastructure & applications

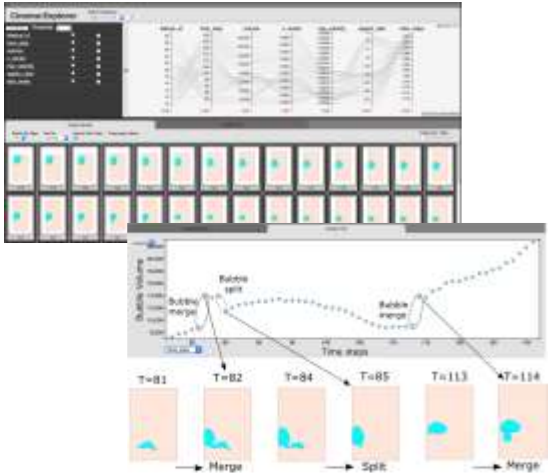

Reduced data abstracts produced in situ (**VTK-m**) can be used in a post hoc Cinema-based workflow to study bubble dynamics in a fluidized bed

Deployment of Cinema Workflows in Spack-based Pantheon reproducible pipelines

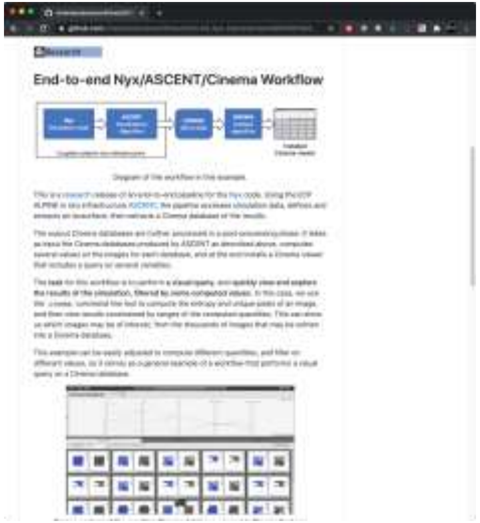

ECP ALPINE Project – built upon long standing ASCR R&D

Summary: ECP's ALPINE team provides data analytics, in situ algorithms, infrastructure and visualization approaches to support scientific discovery and understanding of data produced by exascale science applications.

- LANL algorithm work includes Sampling (Biswas, Hazarika) & Statistical Feature Extraction (Dutta).
- LANL (Ahrens) provides ECP leadership as the Data & Vis L3 and as the ALPINE project lead.
- LANL (Turton) provides project management for ALPINE and the Data & Vis portfolio.

Next Steps: ALPINE is focusing on integration into ECP science applications and with other ECP data & visualization software technologies including Cinema, VTK-m, HDF5.

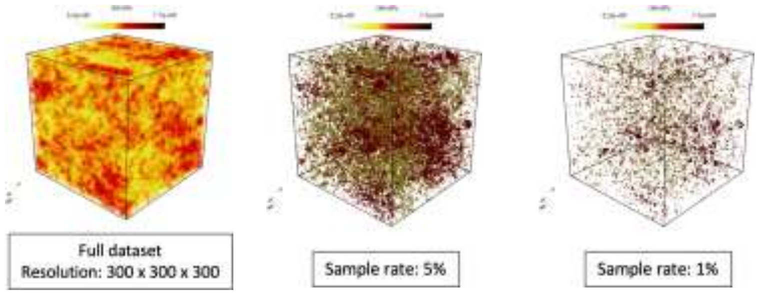

Sampling is applied to the ECP Nyx cosmology code, preserving dark matter halo features. Data-driven sampling preferentially selects and saves more important data while reducing I/O.

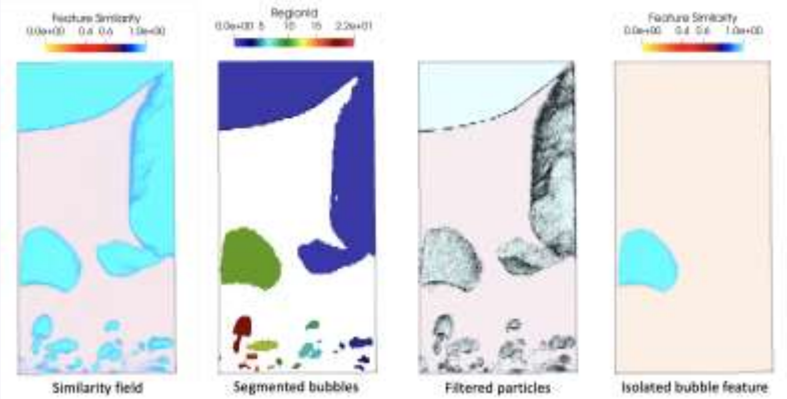

Statistical feature extraction identifies bubbles in situ in an ECP MFIX-Exa. This workflow reduces I/O by two orders of magnitude while enabling greater temporal fidelity.

Molecular Information Storage (MIST): ADS Codex

Advantages and Challenges

- Advantages
  - Extreme data density (33 zettabytes can be stored into a ping pong ball)
  - Storage longevity
  - Cheap copies (PCR)
- Challenges
  - High error rate (≈1%)
  - New error types (insertions and deletions)
  - Stochastic “quality” scores
  - Extreme read/write asymmetry

Goals and Design

- Goals
  - Adapt to quickly evolving technologies
  - Provide very high data density
  - Available to “brave users” in 5-8 years
- Features
  - Metadata based assembly
  - Novel algorithms that can adapt to the systematic errors
  - Optimized decoding scheme that can recover from insertion and deletion errors

ASCR LANL Reverse Site-Visit Quantum Information Science

Andrew Sornborger

12 May 2022

LANL has Developed a Quantum Information Science Strategy as Part of Broader 10 Year Computing Goals

Information Science and Technology Pillar Goals:

- Execute applications of mission relevance on large, fault-tolerant, modular quantum computers.
- Be recognized as a major influence in quantum algorithm development across the entire software stack for both fault-tolerant quantum computers and NISQ devices.

The main organizing principle among our quantum researchers is the pursuit of quantum advantage (i.e., a provable speedup over classical for a *useful* quantum algorithm) by incorporating our expertise's in efficient quantum algorithms.

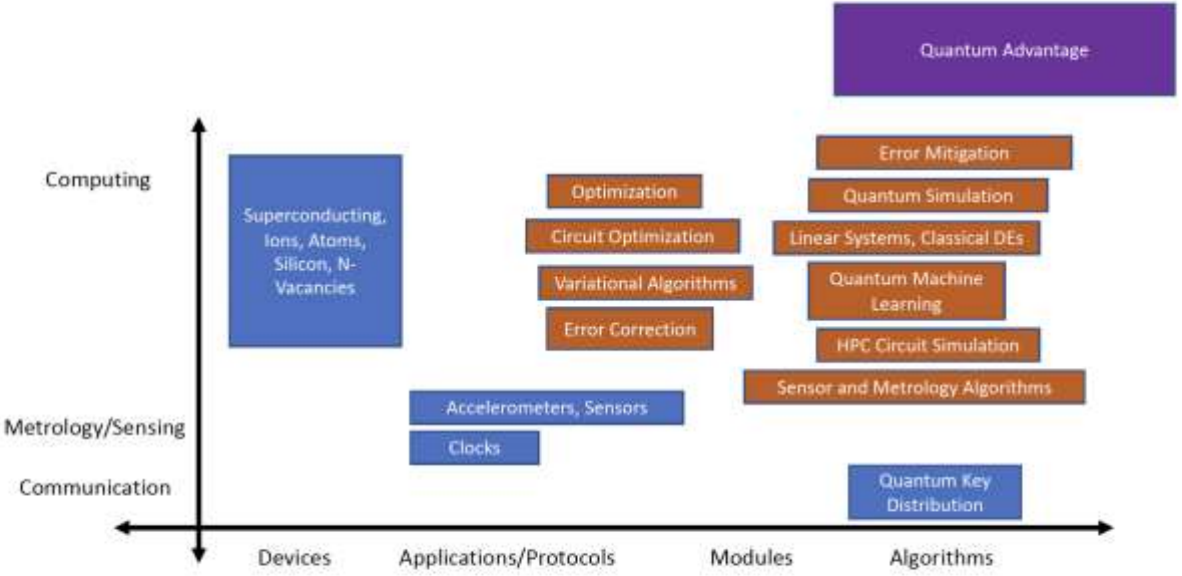

LANL Funding for QIS ($14.4M)

Funding agencies and contributions:

- ASCR ($5,445K)
  - Fundamental Algorithmic Research for Quantum Computing (Somma, $241K)
  - Optimization, Verification, and Engineered Reliability of Quantum Computers (Coles, $269K)
  - Advancing Integrated Development Environments for Quantum Computing through Fundamental Research (Coles, $285K)
  - Quantum Science Center (Sornborger, $4,650K)

- HEP QuantISED ($1.2M)
  - Quantum Field Theories as Spin Models on Quantum Computers (Bhattacharya)
  - Quantum Computing for Neutrino-Nucleus Dynamics (Gupta)
  - Quantum Machine Learning for Lattice QCD (Yoon)
  - Quantum Foundations on Quantum Computers (Sornborger, Albrecht (UC Davis))
  - Renormalization of Entanglement in Quantum Field Theories (Alves)

- ASC ($2.4M)
  - Beyond Moore's Law quantum program

LANL Funding for QIS ($14.4M)

Funding agencies and contributions:

- LDRD DR ($4,738K)
  - Machine Learning for Realizing Next-Generation Quantum Hardware $1,663K (Martin, current)
  - Quantum Chemistry using Quantum Computers $1,600K (Dub, ending)
  - Taming Defects in Quantum Computers $1,475K (Paking, ended)
  - Invited Full Proposals – possible funding FY2023
    - Simulation of Materials on Quantum Computers
    - Quantum Advantage for Materials Science
    - Accelerating Scientific Discovery with Quantum Annealing
    - Next-Generation Quantum Sensors

- LDRD ER ($325K)
  - Digital Quantum Simulations for Dynamical Quantum Matter $325K (Zhu, current)
  - Invited Full Proposals – possible funding FY2023
    - Entangled Two Photon Microscopy (Werner)
    - Quantum Enhanced Atom Interferometer (Ryu)
    - Developing Novel Techniques to Detect and Manipulate Quantum Spin Liquids (Lee)

- LDRD DI ($320K)
  - Quantum Algorithm Development for Optimization $320K (Eidenbenz, ended)

Quantum Information Science ASCR Portfolio, $5.5M

Funding up from $3.7M in FY21

- FAR-QC: Fundamental Algorithmic Research for Quantum Computing (Lead: Ojas Parekh, SNL. LANL PI: Rolando Somma) $341K
- OVER-QC: Optimization, Verification, and Engineered Reliability of Quantum Computers (Lead: SNL, LANL PI: Patrick Coles) $269K
- AIDE-QC: Advancing Integrated Development Environments for quantum computing Through Fundamental Research (Multi-lab: LANL PI: Patrick Coles) $285K
- The Quantum Science Center: A DOE Quantum Information Science Research Center (Lead: ORNL, LANL PI: Andrew Sornborger) $4,650K)

Patrick Coles

Rolando Somma

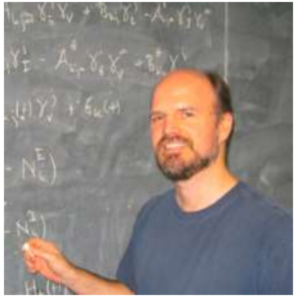

Andrew Sornborger

FAR-QC: Fundamental Algorithmic Research for QC

Lead: SNL. LANL PI: Rolando Somma - $341K

FAR-QC seeks to deliver quantum algorithms that offer provable asymptotic advantages over the best-known or best-possible classical counterparts

Rolando Somma
Theory and Practice
Interface Lead

Yigit Subasi

Key Personnel

Recent Scientific Achievement: Verifying the solution to the quantum linear systems problem is as hard as solving it

Impact: New methods for the quantum linear systems problem, based on variational and related optimization approaches, can be shown to be less efficient.

PRX QUANTUM 2, 010315 (2021)

Complexity of Quantum State Verification in the Quantum Linear Systems Problem

Rolando D. Somma[1,*] and Yiğit Subaşı[2,†]

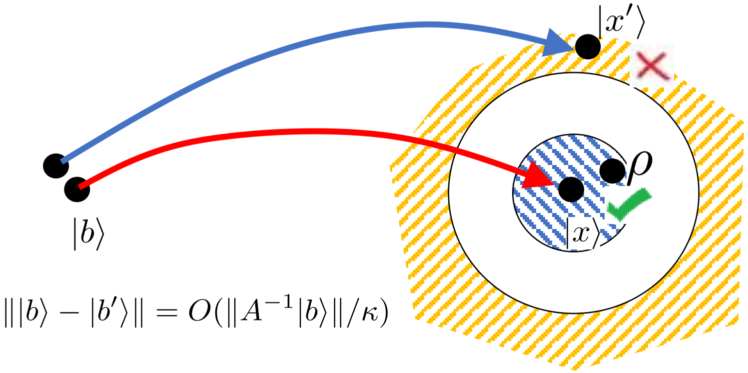

OVER-QC: Optimization, Verification, and Engineered Reliability of Quantum Computers

Lead: SNL. LANL PI: Patrick Coles, $269K

Development of methods for mitigating errors on quantum computers e.g., via machine learning or using classically simulable circuits.

Analysis of the impact of noise on near-term quantum algorithms. Discovered a new phenomenon of noise-induced flat landscapes in variational quantum algorithms.

Patrick Coles

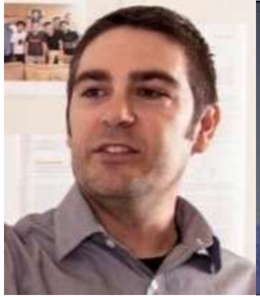

Lukasz Cincio

Marco Cerezo

Key Personnel

Recent Scientific Achievement: Variational Quantum State Eigensolver algorithm for extracting the eigenvalues and eigenvalues of a density matrix, with application to error mitigation, entanglement spectroscopy, and principal component analysis (PCA).

Impact: Makes large-scale PCA a more near-term quantum application

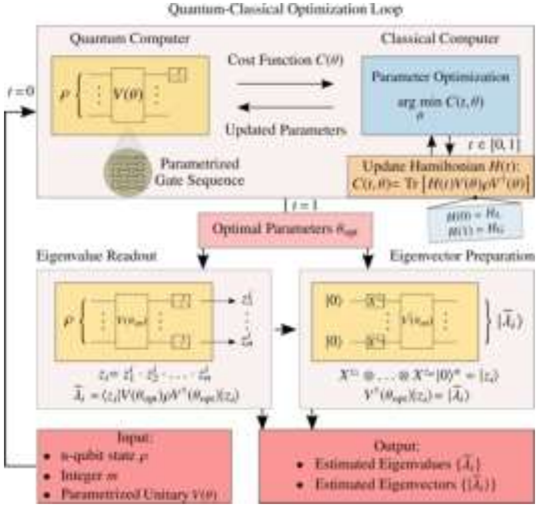

AIDE-QC: Advancing Integrated Development Environments for quantum computing Through Fundamental Research

Lead: LBNL. LANL PI: Patrick Coles, $285K

Analysis of trainability of quantum neural networks and variational quantum algorithms

Development of methods to improve the trainability of quantum neural networks e.g., correlating parameters and reducing circuit depth

Patrick Coles

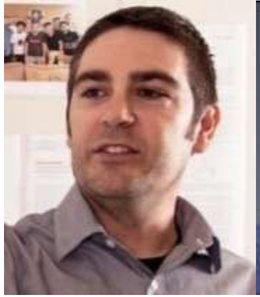

Lukasz Cincio

Yigit Subasi

Key Personnel

Recent Scientific Achievement: Showed that entanglement leads to an apparent violation of the classical No-Free-Lunch Theorem. Developed a quantum No-Free-Lunch (NFL) theorem, where the limit on the learnability of a unitary is reduced by entanglement.

Impact: Establishes that entanglement is a resource for quantum machine learning and allows for efficient scaling of quantum neural networks.

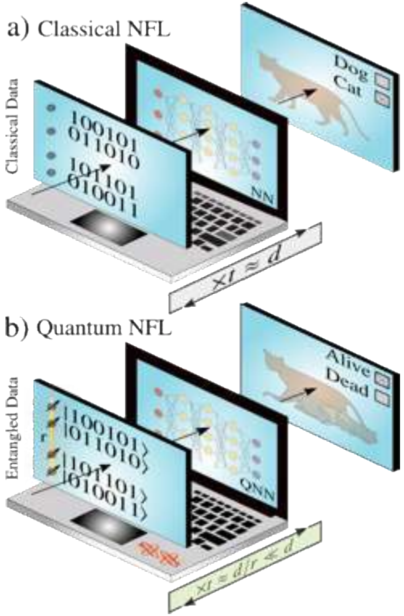

Quantum Science Center (QSC)

Lead: ORNL, LANL PI: Andrew Sornborger, $4,650K (FY22) (FT21 – FY25 $21.4M)

Goal: Overcoming key roadblocks in quantum state resilience, controllability, and ultimately scalability of quantum technologies

LANL is performing research for all three research thrusts

LANL PI: Andrew Sornborger, Quantum algorithms and simulations thrust leader

Thrust 1
Address the fragility of quantum states through the **design of new topological materials** for QIS

Accelerated quantum information processing

Thrust 2
Develop **algorithms and simulation** for computation and sensing with current and future QIS hardware

Prediction of new physical and material behaviors

Thrust 3
Design new **quantum devices and sensors** to detect dark matter and topological quasiparticles

New quantum sensing capabilities to explore the previously unmeasurable

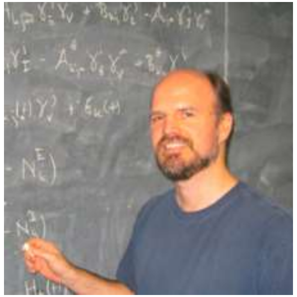

Thrust 2: Quantum Algorithms and Simulation

Thrust 2 will achieve predictive capabilities for the study of strongly coupled quantum systems, including topological systems and quantum field theories, and develop and test quantum algorithms for quantum-limited sensors. To achieve this goal, QSC researchers are developing efficient, scalable, and robust quantum simulation and metrology algorithms, testing these algorithms in predictive dynamical quantum simulation and quantum sensing applications, and developing software tools to support algorithm analysis, optimization, and implementation.

Led by LANL's Andrew Sornborger

Quantum Science Center (QSC)

Lead: ORNL, LANL PI: Andrew Sornborger, $4,650K (up from $2.6M in first year)

Key personnel: LANL is involved in 10 QSC research projects.

Rosa, Thomas, Tutchton, Yoo, Zhu - Topological materials prediction, synthesis, materials development

Zapf, Movshovich - Quantum spin liquid materials

Carlson, Gupta, Alves, Neill - Strong interactions and dynamics: from quarks to nuclei

Somma, Subasi - Quantum simulation algorithms that optimally exploit problem structure

Coles, Cincio, Somma, Subasi - Error mitigation on near-term quantum devices

Eidenbenz - Scalable quantum and classical algorithms and software technologies for efficient modelling and simulation of quantum systems

Htoon - Hybrid quantum sensors

Crook - Squeezed readout of quantum sensor

Martin, Boshier, Somma, Coles, Sornborger - Algorithms and implementations for robust quantum metrology and quantum sensing

QSC Spotlight on Management Structure

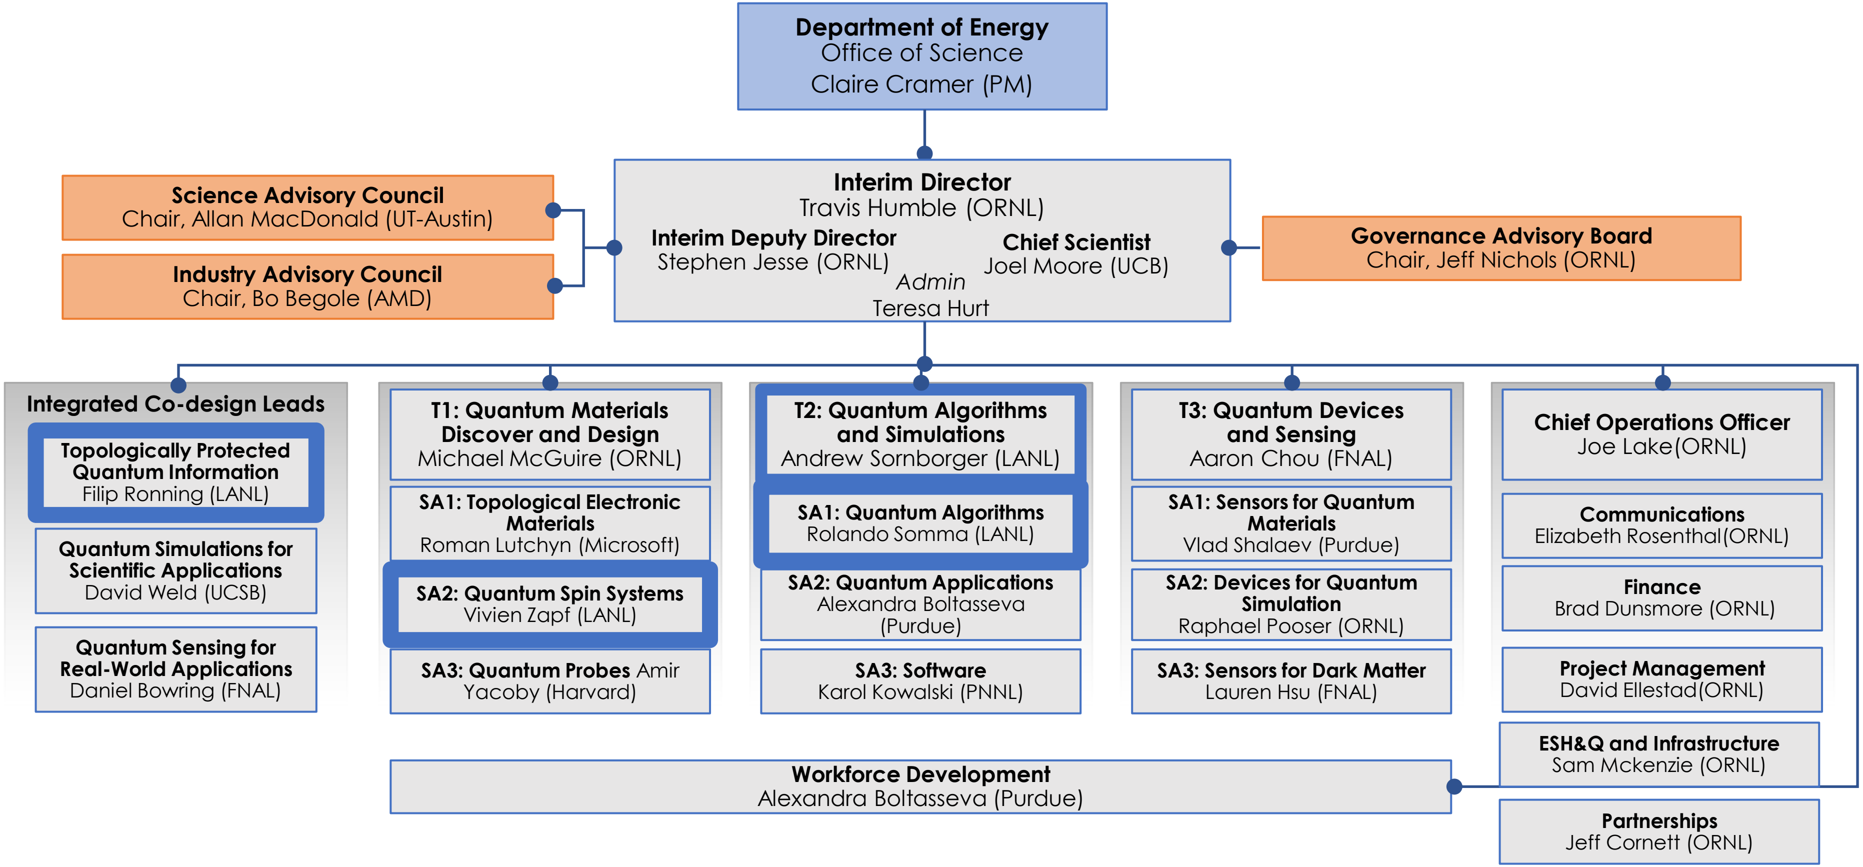

Los Alamos NATIONAL LABORATORY

Quantum Science Center Organization
April 2022

QUANTUM SCIENCE CENTER

Hamiltonian simulation of low-energy quantum dynamics

Scientific Achievement

An improved Hamiltonian simulation method for low-energy states.

Significance and Impact

Simulation of physical systems can be designed more cleverly such that optimal resources are allocated depending on problem specific features, such as energy constraints.

Research Details

— State-of-the-art product formula, e.g., Trotter-Suzuki, methods do not consider physical problem constraints.

This is important for problems where the reliable physics happen in a subspace, such as low-energies.

— Result 1: Exponential suppression of leakage to high energies caused by the product formula.

— Result 2: The true effective-low energy dynamics ~ Dynamics generated by product formula.

Şahinoğlu, Burak, and Rolando D. Somma. "Hamiltonian simulation in the low-energy subspace." ***npj Quantum Information*** 7, no. 1 (2021): 1-5.

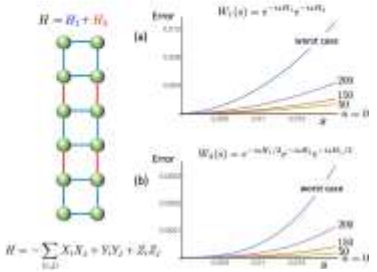

Fig.1:A small scale simulation of errors indicating a more efficient simulation at low-energy.

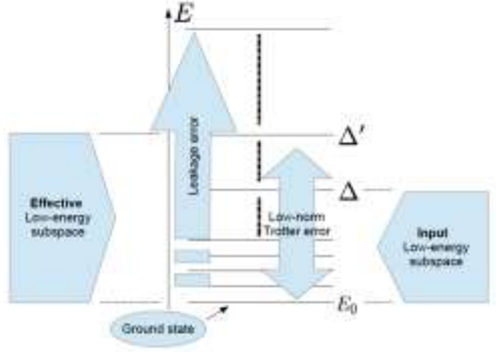

Fig.2: Given an initial state supported in the low-energy subspace, the complexity of product formulas can be improved. To this end, we introduce the notion of effective low-energy Hamiltonians and prove several bounds on the exponential decay of the leakage to subspaces of high energies. Main result follows from balancing low-norm Trotter error and leakage error.

QSC Spotlight on S&T Innovation Chain

Quantum algorithm for learning unitary operations by project 2.2.6 at LANL

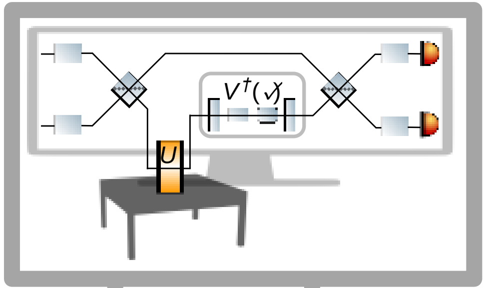

CV

DV

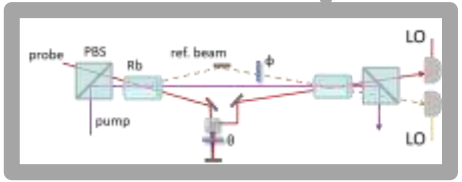

Implementation on quantum optical platform by project 2.2.1 at ORNL

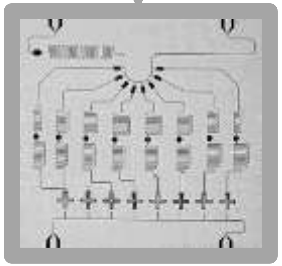

Implementation in transmon platform by project 3.1.2 at Purdue

Progress/Achievement

Cross-thrust efforts launched to explore quantum algorithms on two different quantum hardware platforms

Significance and Impact

A new algorithm drives collaboration to demonstrate the power of quantum platforms to improve learning unitary operations. The co-design opportunity provides feedback on experimental capabilities of continuous variable (CV) and discrete variable (DV) platforms

Details

- LANL provided a recent quantum compiling algorithm for both discrete and continuous variable implementations
- ORNL and Purdue began constructing a demonstration apparatus once it was determined that the state preparation was congruent with ongoing milestone tasks
- Initial sensemaking and problem identification managed by Fermilab are complete. Next trimester we will develop a comprehensive research plan

Volkoff, Tyler, Zoë Holmes, and Andrew Sornborger. "Universal Compiling and (No-) Free-Lunch Theorems for Continuous-Variable Quantum Learning." ***PRX Quantum*** 2, no. 4 (2021): 040327.

Work was performed at ORNL, LANL, Purdue University, and FNAL as part of projects 2.2.1, 2.2.6, 3.1.2 and 4.3.9

LDRD DR: Quantum Chemistry on Quantum Computers

Grand-goal: create a platform (hybrid algorithms) for computational molecular chemistry using NISQ gate-based & quantum annealers platforms

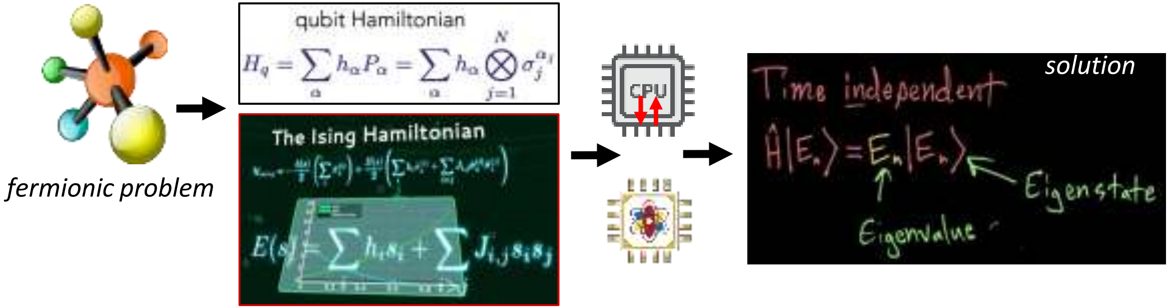

Pavel Dub
Project Lead

Increased depth can both degrade the circuit accuracy and reduce trainability. We propose an approach to reduce ansatz circuit depth. Our approach adds an additional optimization loop to the VQE that permutes qubits in order to solve for the qubit Hamiltonian that maximally localizes correlations in the ground state. For representative molecular systems, LiH, H2, (H2)2, H×=4 , H+3 , and N2, we demonstrate that placing entangled qubits in close proximity leads to shallower depth circuits required to reach a given eigenvalue-eigenvector accuracy.

Tkachenko, Nikolay V., James Sud, Yu Zhang, Sergei Tretiak, Petr M. Anisimov, Andrew T. Arrasmith, Patrick J. Coles, Lukasz Cincio, and Pavel A. Dub. "Correlation-informed permutation of qubits for reducing ansatz depth in the variational quantum eigensolver." ***PRX Quantum* 2**, no. 2 (2021): 020337.

LDRD ER: Quantum Algorithms for Combinatorial Optimization Problems

Evaluating quantum hardware ‘in the Cloud’. Quantum volume is an indicator of both how many qubits and how deep of a quantum circuit may be used effectively on a gate-based quantum computer.

Vendor	Backend	Qubit count	Published QV (vendor)	High optim-ization QV	User-expe-rienced QV	Vendor/User QV (Ratio) “low=good”
IBM Q	Montreal	27	128	32	8	16
	Brooklyn	65	32	32	8	4
	Perth	7	32	8	8	4
	Manila	5	32	16	16	2
Quantinuum	H1-2	12	2048	512*	512*	4*
IonQ	IonQ device	8	-	8	8	-
Rigetti	Aspen-11, Aspen-M-1	38, 79	8	1	1	8
OQC	Lucy	8	-	1	1	-

Stephan Eidenbenz
Project Lead

- Ongoing testing for QV 1024 for Quantinuum looks promising, credit constrained
- 18 out of 19 IBM Q backends had user QV=8
- Intense optimization improved most IBM Q backend QV, but not likely to help most users
- Compilation complexity for ion trap systems (Quantinuum, IonQ) significantly lower, but generally viewed as limited in qubit scale and slow execution

Pelofske, Elijah, Andreas Bärtschi, and Stephan Eidenbenz. "Quantum Volume in Practice: What Users Can Expect from NISQ Devices." *arXiv preprint arXiv:2203.03816* (2022).

ASC Beyond Moore's Law Quantum Portfolio, $2.9M

NNSA ASC BML is working with non-von Neumann computational hardware in order to find areas of computational advantage for these systems relative to 'standard' computational hardware such as CPUs and GPUs.

Recent results

- Unified framework for quantum algorithms for classical, non-linear PDEs. Working with PsiQuantum to estimate quantum resources.
- Quantum on error mitigation and barren plateaus.
- Studied fidelity of Dicke states achievable through circuit optimization and state tomography measurements on IBMQ quantum devices.
- Non-linear transformations in quantum computation

Provable generalization learning quantum simulations for near-term QCs

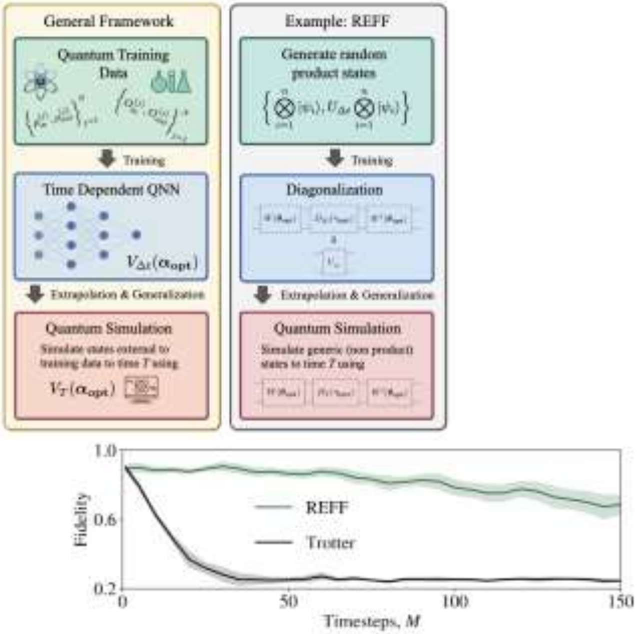

Gibbs, Joe, Zoë Holmes, Matthias C. Caro, Nicholas Ezzell, Hsin-Yuan Huang, Lukasz Cincio, Andrew T. Sornborger, and Patrick J. Coles. "Dynamical simulation via quantum machine learning with provable generalization." *arXiv preprint arXiv:2204.10269* (2022).

Los Alamos NATIONAL LABORATORY

LANL is building a quantum community and workforce with postdoctoral researchers, summer schools

The Quantum Summer School is building future workforce

- 2022 – 5th quantum school – in person participation!
- 2022: 490 applications for 20 positions
- Format: lectures by top scientists (Aaronson, Martinis, McClean,...), projects with LANL mentors (12 mentors)
- 2021 Research Projects:
  - Unified Approach to Data-driven Quantum Error Mitigation
  - Variational Quantum Algorithm for Estimating the Quantum Fisher Information
  - A Generalized Measure of Quantum Fisher Information
  - Correlation-Informed Permutation of Qubits for Reducing Ansatz Depth in VQE
  - Noise-Induced Barren Plateaus in Variational Quantum Algorithms
  - Absence of Barren Plateaus in Quantum Convolutional Neural Networks
  - Dynamical simulation via quantum machine learning with provable generalization
- We are tracking student outcomes (publications, starting to see hiring)

Lukasz Cincio

Patrick Coles

Marco Cerezo

Yigit Subasi

Marco Cerezo,
Early Career
UNLP Argentina

Zoe Holmes
Imperial College London
2019 Quantum School
1st Kac Fellow 2021
EPFL Sandoz Asst. Professor

Andreas Baertschi,
Early Career
ETH Zurich

Cinthia Huerta Alderete
UMD, Instituto Nacional de Astrofísica, Óptica y Electrónica, INAOE, Mexico

Crosscut: Artificial Intelligence/Machine Learning for Science

Aric Hagberg
Computer, Computational, and Statistical Science Division

May 12, 2022

ASCR portfolio projects with ML crosscut

Math

- ASCR base program: ML strategies for compression of particle data for checkpoint restart. SNL, (Chacon)
- **Data Intensive ML: Inertial neural surrogates for stable dynamical prediction LANL (Tang, FY22 start)**

Computer Science

- **Legion: core technology for parallelizing ML LANL (McCormick)**

Partnerships

- **Probabilistic Sea-Level Projections from Ice Sheet and Earth System Models. LANL (Price)**
- **Nuclear Computational Low Energy Initiative (NUCLEI), LANL, (Carlson)**
- **Tokamak Disruption Simulation, LANL (Tang)**
- **FusMatML: Machine Learning Atomistic Modeling for Fusion Materials, LANL, (Perez)**
- High-Fidelity Boundary Plasma Simulation (HBPS): ML strategies for particle collisions in plasmas (Chacon)
- Towards Exascale Astrophysics of Mergers and Supernovae (TEAMS), ORNL, (Fryer)
- Computing the Properties of Matter with Leadership Computing Resources, TJNAL (Yoon)
- Accelerating HEP Science: Inference and Machine Learning at Extreme Scales. ANL (Lawrence)
- Development of Terrestrial Dynamical Cores for the ACME to Simulate Water Cycle, PNNL, (Karra)
- Leveraging ML/AI techniques to enable a breakthrough in ultra-short-bunch, SLAC (Scheinker)

Quantum Science Center

LANL Artificial Intelligence Portfolio

- Developing robust AI capabilities that enable new research and development opportunities for LANL mission applications – not general AI
- Leveraging historical strengths in computational physics – “physics-informed ML” – integration of AI with modeling and simulation
- Investing with LDRD program to develop capability and teams
- Increasing portfolio of DOE SC, LDRD, DOE Applied Energy, NNSA, and Global Security projects
- Access to ML hardware and partnerships with ML hardware vendors
- Engaging community through LANL-led workshops and conferences
- Developing pipeline with structured student internships
- *We are actively responding to ASCR and SC funding opportunities*

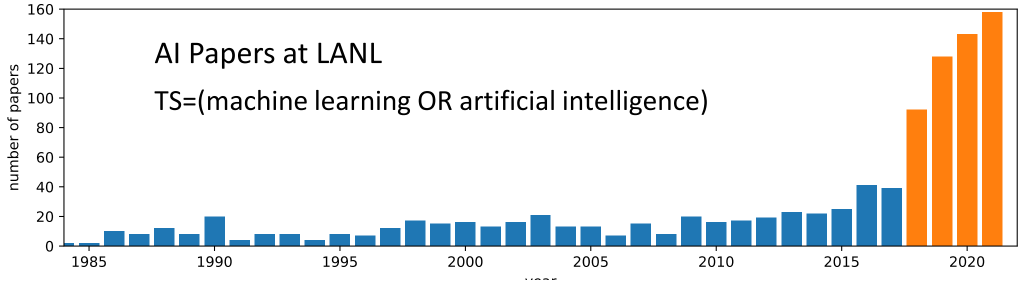

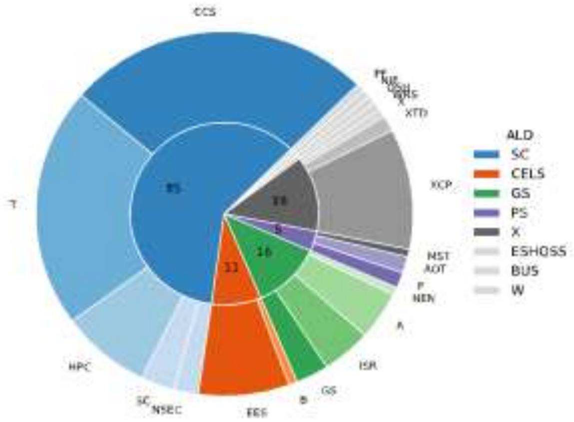

AI workforce spans laboratory > 250

AI Capability Themes at Los Alamos
Connected to Mission Sponsors and Capabilities

Physics-informed Machine Learning (NNSA/ASC, SC, OE, FECM)	AI for Computational Science (NNSA/ASC, SC/BES)
Experimental and Observational Workflows (NNSA/ASC, ECP)	AI for Sensing (Global security, NNSA/NA-22, SC/BER)

Core AI Technologies

Physics-informed ML: The use of scientific knowledge in machine learning techniques to improve learning and predictions. The incorporation of physical constraints and equations when learning from data.

Computational Science: Integration of AI methods with high-performance computing algorithms and simulations. The use of machine learning for scale-bridging and for emulators for high-cost simulations.

Workflow Optimization: Reducing the time to result for computational, experimental, observational, or production workflows. This includes experimental design, dynamic control of workflows, and analysis of outputs.

Sensing: The development and deployment of machine learning algorithms for and with sensors. The design of machine learning inference algorithms for process controls, remote low-power sensors, and manufacturing.

Core Technologies: Mathematics and statistics fundamentals including uncertainty quantification and optimization, software tools, and HPC implementations. Industrial and academic R&D are providing advanced technologies, but gaps exist for their use on our mission problems.

AI Capability Themes at Los Alamos
Crosscut to 2022 AI Workshop Topics

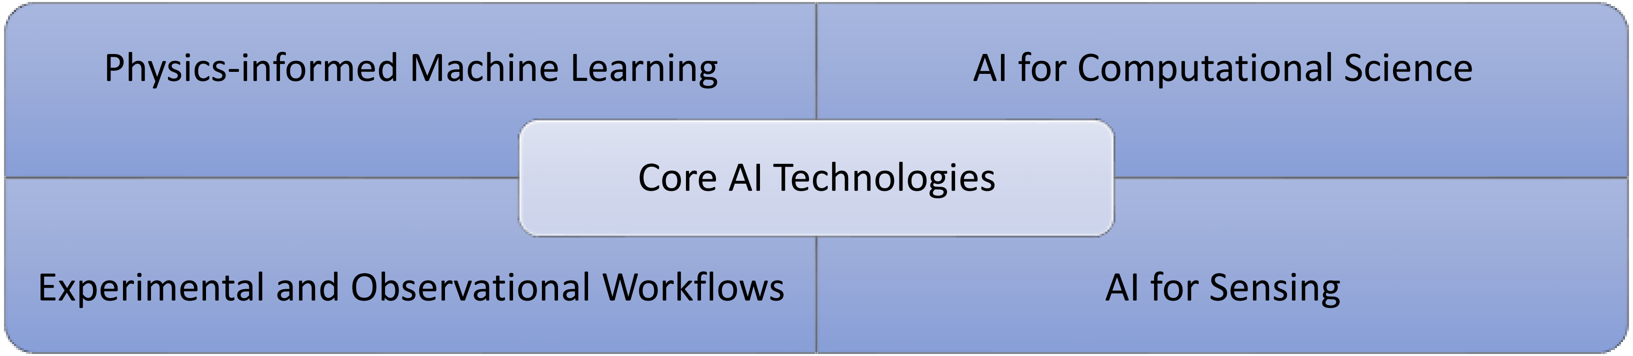

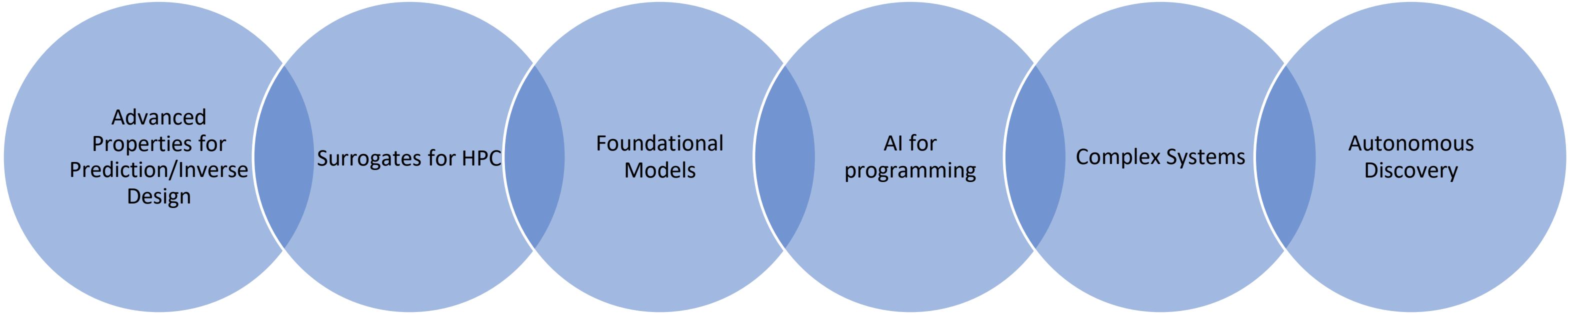

Technology: LDRD Program is Investing in Core AI R&D (~$15M/year)

DOE ASCR Scientific ML Priority Research Directions

Scientific Machine Learning: Foundations

Domain-Aware: Leverages & respects scientific domain knowledge. Physics principles, symmetries, constraints, uncertainties & structure-exploiting models

Interpretable: Explainable and understandable results. Model selection, exploiting structure in high-dimensional data, use of uncertainty quantification with machine learning

Robust: Stable, well-posed & reliable formulations. Probabilistic modeling in ML, quantifying well-posedness, reliable hyperparameter estimation

Scientific Machine Learning: Capabilities

Data-Intensive SciML: Scientific inference & data analysis. ML methods for multimodal data, in situ data analysis & optimally guide data acquisition

Machine Learning-Enhanced Simulations: ML hybrid algorithms & models for predictive scientific computing. ML-enabled adaptive algorithms, parameter tuning & multiscale surrogate models

Intelligent Automation and Decision Support: Adaptivity, automation, resilience, control. Exploration of decision space with ML, ML-based resource management, optimal decisions for complex systems

LANL LDRD Research Projects FY20-

Differentiable Programming: Bridging the Gap between Numerical Models and Machine Learning Models, Dan O'Malley

Topological Relation-Based Image Analysis using Graphs, Diane Oyen

Uncertainty Quantification for Robust Machine Learning, Diane Oyen

Deep Learning in a Noisy World: Algorithms for robust training and predictive uncertainty, Sunil Thulasidasan

Learn Equations, not Models: Interpretable Digital Twins for Complex Systems, Arvind Mohan

The Optimization of Machine Learning: Imposing Requirements on Artificial Intelligence, Russell Bent

Warm-starting Quantum Machine Learning, Lukasz Cincio

Machine Learning for Realizing Next-Generation Quantum Hardware, Michael Martin

Tensor Networks: Robust Unsupervised Machine Learning for Big-Data Analytics, Boian Alexandrov

Sampling the unknown: Robust modeling of atomic potentials, Kipton Barros

In-Situ Inference: Bringing Advanced Data Science into Exascale Simulations, Earl Lawrence

Enabling Predictive Scale-Bridging Simulations through Active Learning, Tim German

Machine Learning for Turbulence, Daniel Livescu

Adaptive Machine Learning for Closely Spaced Ultra-Short Intense Accelerator Beams , Alex Scheinker

AI@Sensing, Satish Karra

Physics-informed Machine Learning

Anomaly Detection in Energy Systems
OE,LDRD, Bent

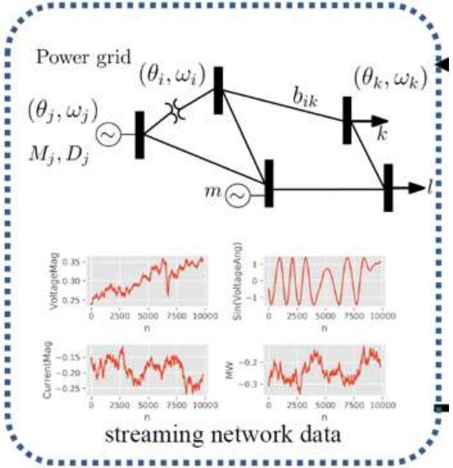

Incorporate known physics of electrical circuits

Machine Learning for Turbulence
LDRD (Livescu)

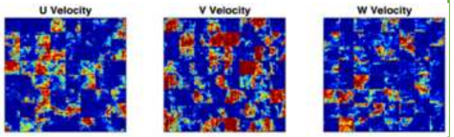

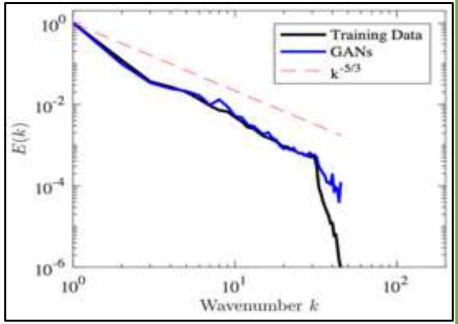

Add physics constraints to black-box ML for turbulence emulators

Differentiable Programming: Bridging the Gap between Numerical and Machine Learning Models
LDRD (O'Malley)

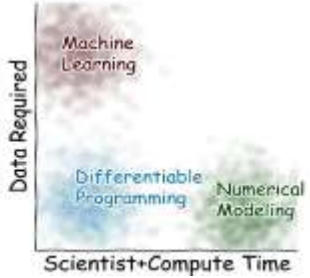

Combine trustworthy numerical modeling with ML to produce fast models

AI for Computational Science

Enabling Predictive Scale-Bridging Simulations through Active Learning

LDRD (Germann)

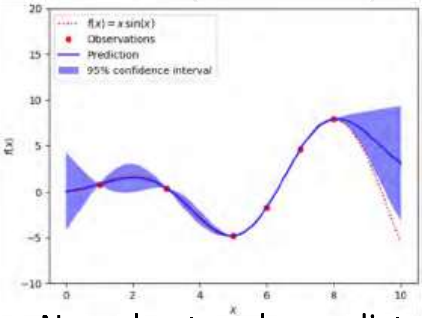

Neural networks predict upscaled parameters and uncertainties in multiscale systems. Materials damage.

Quantum Chemistry and Molecular Dynamics

LDRD, BES, ASC (Barros)

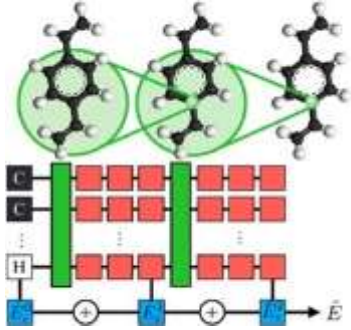

ML constructs potentials based upon large datasets of quantum calculations. Materials properties, aging.

In-situ Inference

LDRD (Lawrence)

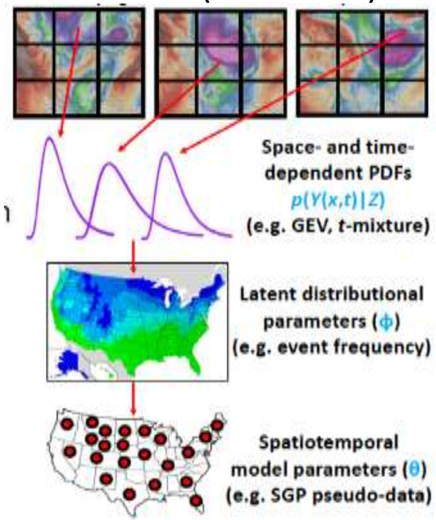

Bringing advanced data science into exascale simulations. Climate, space weather.

AI for Sensing

Earthquakes
BES, LDRD
(Johnson)

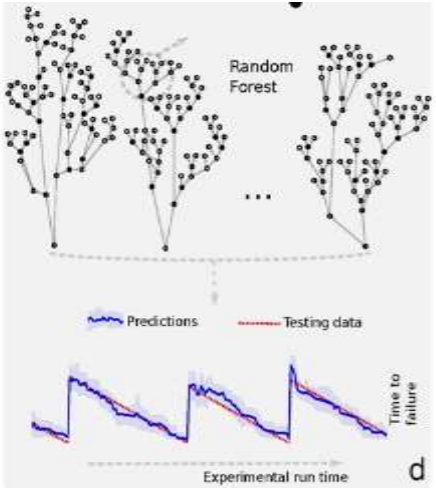

ML to probe fault physics and assess earthquake hazard

Structural Health
Monitoring
LDRD (Farrar)

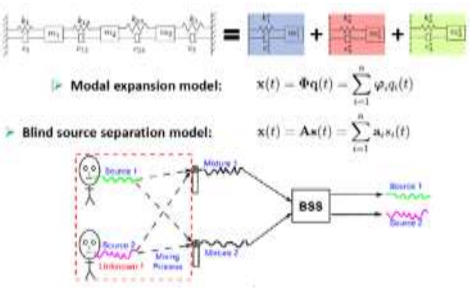

Remote measurement of high-resolution full-field spatio-temporal dynamics

Manufacturing
NNSA (Wachtor)

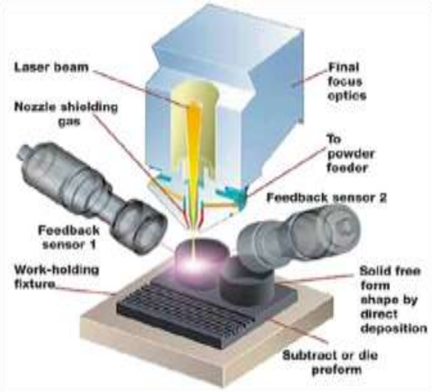

Monitor formation of pores using acoustic emissions data + machine learning methods

Experimental and Observational Workflows

Reinforcement Learning for Scientific Environments
ECP (Sweeney)

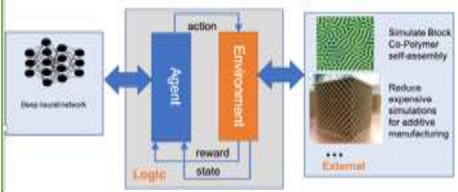

ExaRL:" Developed as part of the Exascale Computing Project ExaLearn

Real-time Adaptive Acceleration of Dynamic Experimental Science
LDRD (Ahrens)

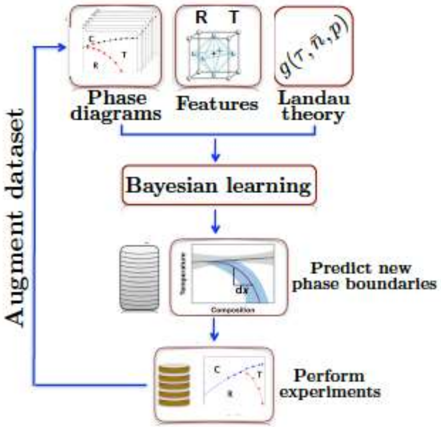

Adaptive Machine Learning for Closely Spaced Ultra-Short Intense Accelerator Beams
LDRD, HEP (Scheinker)

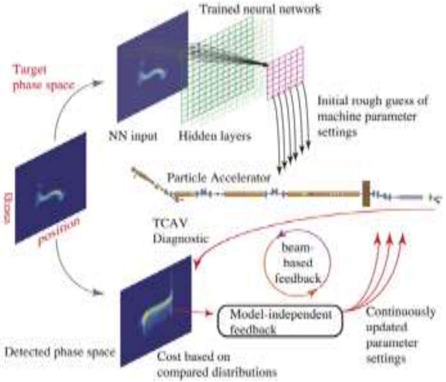

Core AI Technologies

Optimization of ML
Constrained, robust, tractable, learning methods
LDRD (Bent)

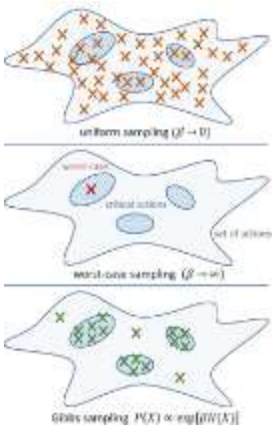

Anomaly detection in energy systems

Graphical models
Probabilistic machine learning methods
LDRD (Lokhov)

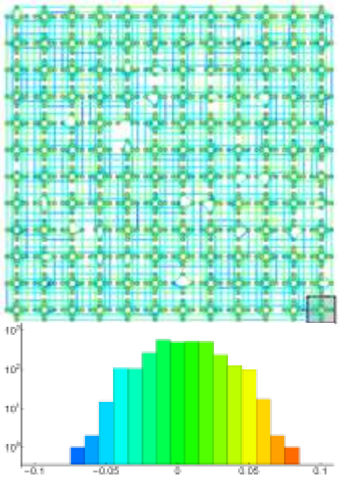

Calibrating the D-Wave quantum annealer

Tensor factorization
Extraction of explainable hidden features in data
LDRD (Alexandrov)

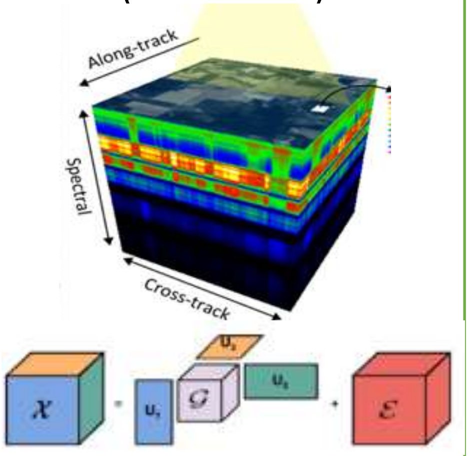

Applications in cancer, earthquakes, ~100TB data

BES: AI-enhanced quantum chemistry, a breakthrough capability (Barros)

Multimodal approaches for leveraging domain knowledge with state-of-the-art machine learning to engineer biocatalysts

Atomic geometry

HIPNN model

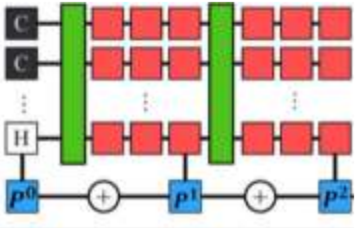

Semi-empirical quantum mechanical model (SEQM)

Corrections to: MNDO, AM1, PM3, ...

Self-consistent predictions (energies, excited states, dynamics, polaritonics)

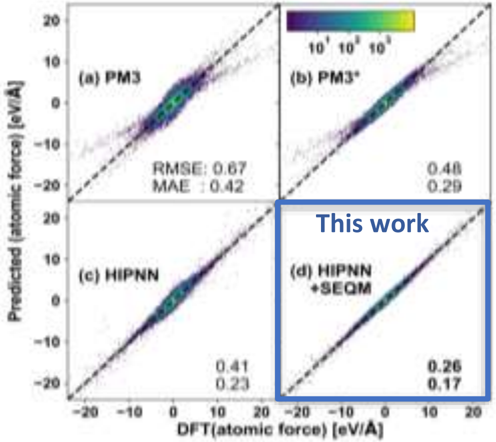

Record performance on the tripeptide/COMP6 benchmark

Scientific Achievement (To appear in PNAS)

We developed a deep learning model which leverages physical insight to achieve a stunning combination of computational efficiency and accuracy.

Significance and Impact

The new ML approach produces state-of-the-art performance on many existing benchmarks, and brings new capabilities, such as the accurate prediction of excited state dynamics and polaritonics.

Research details

Our next generation ML models are trained to predict a corrections to a semi-empirical Hamiltonian. This reduced-order quantum model is solved self-consistently and compared to reference DFT data. Training involves backpropagation through the full work-flow.

U.S. DEPARTMENT OF ENERGY | Office of Science

Los Alamos NATIONAL LABORATORY

Partnerships: Advanced ML Computing Resources

- SambaNova ML hardware and tools
  - Showing ~3x speed up over Nvidia GPU
- Groq-ML accelerators
- LANL partnership with NVIDIA will bring in a large CPU-GPU system in 2023 – Venado.

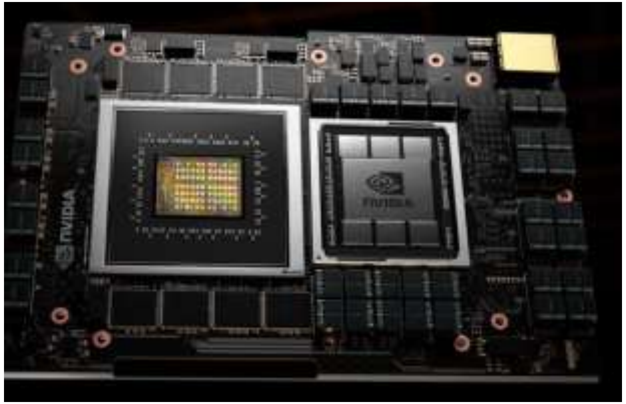

Partnerships: NSF Center AI4Opt

AI Institute for Advances in Optimization

Georgia Tech

- "delivering a paradigm shift in automated decision making at massive scales by **fusing AI and Mathematical Optimization** (MO), to achieve breakthroughs that neither field can achieve independently"
- Led by Georgia Tech – 6 universities, 8 industrial/lab partners
- Synergistic with LANL LDRD DR project "The Optimization of Machine Learning: Imposing Requirements on Artificial Intelligence" (Russell Bent):
  - organize joint workshops/events with the institute,
    - Kickoff Feb 3, 2022
    - Winter school @Santa Fe 2023
  - host graduate students and postdocs from the institute (2 summer 2022)
  - new research collaborations

Community: Convening Workshops and Conferences

- Physics Informed Machine Learning
  - 2016,2018,2020,2022
- Machine Learning in Solid Earth Geoscience
  - 2018,2019,~~2020~~, 2022
- Conference on Data Analysis
  - 2012,2014,2016,2018,2020, ~~2022,~~ 2023
- Machine Learning for Computational Fluid and Solid Dynamics
  - 2019, ongoing international seminars

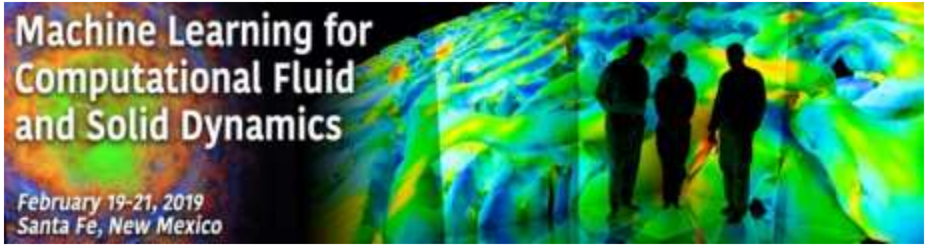

Pipeline: Machine Learning

Applied Machine Learning Summer Research Internship

- Collective advertising helps recruit a large applicant pool with highly-qualified students
- Students come from various academic disciplines - mostly not computer science
- Team Science: Students work on our research projects in small teams with teams of mentors – provides typical lab team science experience

AML 2021 - Virtual

6th year of program

- 2017 Class 12 students
- 2018 Class 24 students
- 2019 Class 15 students
- 2020 Class 16 students
- 2021 Class 12 students

Typically, about 10x applicants.

We track outcomes and have former students who have returned as graduate students, postdocs, and staff